

FortiADC Handbook

VERSION 4.7.4

FORTINET DOCUMENT LIBRARY

<http://docs.fortinet.com>

FORTINET VIDEO GUIDE

<http://video.fortinet.com>

FORTINET BLOG

<https://blog.fortinet.com>

CUSTOMER SERVICE & SUPPORT

<https://support.fortinet.com>

FORTIGATE COOKBOOK

<http://cookbook.fortinet.com>

FORTINET TRAINING SERVICES

<http://www.fortinet.com/training>

FORTIGUARD CENTER

<http://www.fortiguard.com>

END USER LICENSE AGREEMENT

<http://www.fortinet.com/doc/legal/EULA.pdf>

FEEDBACK

Email: techdocs@fortinet.com



Wednesday, October 11, 2017

FortiADC Handbook 4.7.4

First Edition

TABLE OF CONTENTS

Change Log	13
Introduction	14
Features	14
Basic network topology	14
Scope	15
Chapter 1: What's New	17
FortiADC 4.7.4	17
FortiADC 4.7.3	17
FortiADC 4.7.2	17
FortiADC 4.7.1	17
FortiADC 4.7.0	18
FortiADC 4.6.2	19
FortiADC 4.6.1	19
FortiADC 4.6.0	19
FortiADC 4.5.3	21
FortiADC 4.5.2	21
FortiADC 4.5.1	21
FortiADC 4.5.0	22
FortiADC 4.4.0	23
FortiADC 4.3.1	24
FortiADC 4.3.1	25
FortiADC 4.3.0	25
FortiADC 4.2.3	26
FortiADC 4.2.1	26
FortiADC 4.2.0	27
FortiADC 4.1	27
FortiADC 4.0 Patch 2	27
FortiADC 4.0 Patch 1	27
FortiADC 4.0	28
FortiADC 3.2.0	28
FortiADC 3.1.0	28
FortiADC 3.0.0	29
FortiADC 2.1.0	29
Chapter 2: Key Concepts and Features	30

Server load balancing.....	30
Feature Summary.....	30
Authentication.....	31
Caching.....	31
Compression.....	32
Decompression.....	32
Content rewriting.....	32
Content routing.....	32
Scripting.....	32
SSL transactions.....	32
Link load balancing.....	33
Global load balancing.....	33
Security.....	33
High availability.....	33
Virtual domains.....	34
Chapter 3: Getting Started.....	35
Step 1: Install the appliance.....	35
Step 2: Configure the management interface.....	36
Step 3: Configure basic network settings.....	39
Step 4: Test connectivity to destination servers.....	43
Step 5: Complete product registration, licensing, and upgrades.....	43
Step 6: Configure a basic server load balancing policy.....	45
Step 7: Test the deployment.....	48
Step 8: Back up the configuration.....	51
Chapter 4: Server Load Balancing.....	53
Server load balancing basics.....	53
Server load balancing configuration overview.....	56
Configuring real server SSL profiles.....	57
Configuring MySQL profiles.....	61
Single-master mode.....	61
Sharding mode.....	62
Creating a MySQL profile.....	65
Creating a MySQL configuration object.....	65
Specifying the MySQL user account.....	66
Configuring MySQL rules.....	67
Configuring sharding.....	67
Configuring client SSL profiles.....	69
Using real server pools.....	73
Configuring real server pools.....	73
Example: Using port ranges and the port 0 configuration.....	78
Configuring persistence rules.....	80
Configuring content routes.....	86

Using content rewriting rules.....	88
Overview.....	88
Configuring content rewriting rules.....	89
Example: Redirecting HTTP to HTTPS.....	91
Example: Rewriting the HTTP response when using content routing.....	99
Example: Rewriting the HTTP request and response to mask application details.....	101
Example: Rewriting the HTTP request to harmonize port numbers.....	103
Compression and decompression.....	104
Configuring compression rules.....	105
Configuring decompression rules.....	107
Using caching features.....	109
Static caching.....	109
Dynamic caching.....	111
Configuring caching rules.....	111
Configuring certificate caching.....	113
Configuring a certificate caching object.....	113
Configuring Application profiles.....	113
Configuring error pages.....	138
Using source pools.....	138
Configuring source pools.....	139
Example: DNAT.....	141
Example: full NAT.....	142
Example: NAT46 (Layer 4 virtual servers).....	143
Example: NAT64 (Layer 4 virtual servers).....	145
Example: NAT46 (Layer 7 virtual servers).....	147
Example: NAT64 (Layer 7 virtual servers).....	149
Configuring authentication policies.....	150
Configuring load-balancing (LB) methods.....	152
Configuring an L2 exception list.....	154
Creating a Web Filter Profile configuration.....	155
Using the Web Category tab.....	156
Configuring virtual servers.....	156
Two Options for virtual server configuration.....	156
Basic virtual server configuration.....	157
Advanced virtual server configuration.....	159
TCP multiplexing.....	164
Using scripts.....	166
Create a script object.....	167
Import a script.....	167
Export a script.....	167
Delete a script.....	167
Chapter 5: Link Load Balancing.....	168

Link load balancing basics.....	168
Using link groups.....	168
Using virtual tunnels.....	169
Link load balancing configuration overview.....	171
Configuring gateway links.....	173
Configuring persistence rules.....	174
Configuring proximity route settings.....	176
Configuring a link group.....	178
Configuring a virtual tunnel group.....	180
Configuring link policies.....	182
Chapter 6: Global Load Balancing.....	184
Global load balancing basics.....	184
Global load balancing configuration overview.....	186
Configuring servers.....	188
Configuring a global load balance link.....	191
Configuring data centers.....	192
Configuring hosts.....	193
Configuring virtual server pools.....	195
Configuring Topologies.....	197
Configuring dynamic proximity.....	197
Configuring persistence.....	198
Configuring an address group.....	199
Configuring remote DNS servers.....	200
Configuring the DSSET list.....	201
Configuring DNS zones.....	202
Configuring DNS64.....	206
Configuring the response rate limit.....	207
Configuring a Global DNS policy.....	208
Configuring general settings.....	209
Configuring the trust anchor key.....	210
Chapter 7: Network Security.....	212
Security features basics.....	212
Managing IP Reputation policy settings.....	212
Configure IP reputation exception.....	214
Using the Geo IP block list.....	215
Using the Geo IP whitelist.....	216
Enabling denial of service protection.....	217
Configuring a firewall policy.....	218
Configuring the firewall connection limit.....	219
Chapter 8: Web Application Firewall.....	221
Web application firewall basics.....	221
Web application firewall configuration overview.....	222

Predefined configuration elements	223
Severity.....	223
Exceptions	223
Configuring a WAF Profile.....	223
Configuring a Web Attack Signature policy.....	225
Configuring a URL Protection policy.....	231
Configuring an HTTP Protocol Constraint policy.....	232
Configuring an SQL/XSS Injection Detection policy.....	236
Configuring WAF Exception objects.....	239
Configuring a Bot Detection policy.....	239
Chapter 9: Authentication Management	242
Configuring user groups.....	242
Using the local authentication server.....	244
Using an LDAP authentication server.....	245
LDAP bind messages.....	245
Simple bind.....	245
Anonymous bind.....	245
Regular bind.....	246
LDAP over SSL (LDAPS) and StartTLS.....	247
Configuring LDAP binding.....	247
Using a RADIUS authentication server.....	249
Using Kerberos Authentication Relay.....	249
Authentication Workflow.....	250
Step 1: Client authentication.....	250
Step 2: Client service authorization.....	250
Step 3: Client service request	251
FortiADC Kerberos authentication implementation.....	251
Configure Authentication Relay (Kerberos).....	251
Using HTTP Basic SSO.....	252
Configure HTTP Basic SSO.....	253
SAML and SSO.....	254
Configure a SAML service provider.....	254
Import IDP Metadata	256
Chapter 10: Shared Resources.....	257
Configuring health checks.....	257
Monitoring health check status.....	265
Creating schedule groups.....	266
Creating IPv4 address objects.....	267
Configuring IPv4 address groups.....	268
Creating IPv6 address objects.....	268
Configuring IPv6 address groups.....	269
Managing ISP address books.....	270

Create an ISP address book object	272
Creating service objects	273
Creating service groups	274
Chapter 11: Basic Networking	276
Configuring network interfaces	276
Using physical interfaces	276
Using VLAN interfaces	277
Using aggregate interfaces	277
Configuring network interfaces	278
Configuring static routes	283
Configuring policy routes	284
Chapter 12: System Management	286
Configuring basic system settings	286
Configuring system time	287
Configuring an SMTP mail server	289
Configuring FortiGuard service settings	289
Pushing/pulling configurations	291
Backing up and restoring the configuration	292
Updating firmware	293
Upgrade considerations	294
Updating firmware using the web UI	294
Updating firmware using the CLI	296
Rebooting, resetting, and shutting down the system	297
Create a traffic group	298
Create a traffic group via the command line interface	299
Create a traffic group from the Web GUI	299
Create administrator users	300
Configure access profiles	302
Enable password policies	305
Configuring SNMP	306
Download SNMP MIBs	307
Configure SNMP threshold	307
Configure SNMP v1/v2	308
Configure SNMP v3	310
Manage and validate certificates	311
Overview	312
Certificates and their domains	312
Prerequisite tasks	313
Manage certificates	314
Generating a certificate signing request	314
Importing local certificates	317
Creating a local certificate group	318

Importing intermediate CAs	319
Creating an intermediate CA group	320
Validating certificates	321
Configure a certificate verification object	321
Importing CRLs	324
Adding OCSPs	325
Importing OCSP signing certificates	327
Importing CAs	328
Creating a CA group	329
HSM Integration	330
Integrating FortiADC with SafeNet Network HSM	330
Preparing the HSM appliance	331
Generating a certificate-signing request on FortiADC	333
Downloading and uploading the certificate request (.csr) file	335
Uploading the server certificate to FortiADC	336
Chapter 13: Logging and Reporting	338
Using the event log	338
Using the security log	345
Using the traffic log	350
Using the script log	359
Using the aggregate log	359
Configuring local log settings	360
Configuring syslog settings	362
Configuring high speed logging	363
Enabling real-time statistics	364
Configuring alert email settings	365
Configuring an alert email recipient	366
Configuring reports	366
Configuring Report Queries	367
Configuring fast reports	370
Viewing reports	371
Viewing the Overall report	372
Viewing the Server Load Balance report	373
Viewing the Link Load Balance report	375
Viewing the Global Load Balance report	376
Viewing the Security report	377
Display logs via CLI	378
Chapter 14: High Availability Deployments	379
HA feature overview	379
HA system requirements	383
HA configuration synchronization	384
Configuring HA settings	385

Monitoring an HA cluster.....	391
Updating firmware for an HA cluster.....	392
Deploying an active-passive cluster.....	393
Overview.....	394
Basic steps.....	396
Best practice tips.....	396
Deploying an active-active cluster.....	396
Configuration overview.....	397
Basic steps.....	398
Expected behavior.....	399
Traffic to TCP virtual servers.....	399
Traffic to HTTP virtual servers.....	403
FTP traffic and traffic processed by firewall rules.....	405
Best practice tips.....	408
Advantages of HA Active-Active-VRRP.....	408
Deploying an active-active-VRRP cluster.....	408
Configuration overview.....	409
Basic steps.....	410
Best practice tips.....	411
Chapter 15: Virtual Domains.....	413
Virtual domain basics.....	413
Enabling the virtual domain feature.....	413
Creating virtual domains.....	414
Assigning network interfaces and admin users to VDOMs.....	414
Virtual domain policies.....	415
Disabling virtual domains.....	416
Chapter 16: SSL Transactions.....	417
SSL offloading.....	417
SSL decryption by forward proxy.....	419
Layer 7 deployments.....	419
Layer 2 deployments.....	421
Profile configurations.....	422
Certificate guidelines.....	425
SSL/TLS versions and cipher suites.....	425
Exceptions list.....	429
SSL traffic mirroring.....	429
Chapter 17: Advanced Networking.....	431
NAT.....	431
Configure source NAT.....	431
Configure 1-to-1 NAT.....	434
QoS.....	436
Configuring a QoS queue.....	437

Configuring the QoS filter.....	437
Configuring the QoS IPv6 filter.....	438
ISP routes.....	439
BGP.....	440
How BGP works.....	440
IBGP vs. EBGP.....	440
Access list vs. prefix list.....	444
Configuring an Access List.....	445
Configuring an Access IPv6 List.....	445
Configuring a Prefix List.....	446
Configuring an IPv6 prefix list.....	447
OSPF.....	447
Reverse path route caching.....	451
Packet capture.....	453
Chapter 18: Best Practices and Fine Tuning.....	455
Regular backups.....	455
Security.....	455
Topology.....	456
Administrator access.....	456
Performance tips.....	457
System performance.....	457
Reducing the impact of logging on performance.....	457
Reducing the impact of reports on system performance.....	457
Reducing the impact of packet capture on system performance.....	457
High availability.....	458
Chapter 19: Troubleshooting.....	459
Logs.....	459
Tools.....	459
execute commands.....	459
diagnose commands.....	460
System dump.....	461
Packet capture.....	462
Diff.....	463
Solutions by issue type.....	464
Login issues.....	464
Connectivity issues.....	465
Checking hardware connections.....	465
Checking routing.....	465
Examining the routing table.....	469
Examining server daemons.....	469
Checking port assignments.....	469
Performing a packet trace.....	469

Checking the SSL/TLS handshake & encryption.....	470
Resource issues.....	470
Monitoring traffic load.....	470
DoS attacks.....	471
Resetting the configuration.....	471
Restoring firmware (“clean install”).....	471
Additional resources.....	474
Chapter 20: System Dashboard.....	475
Status.....	478
Data Analytics.....	479
Server load balance.....	480
Select a display option.....	480
Filter virtual servers onscreen.....	481
Add virtual servers.....	482
Link load balance.....	482
Global load balance.....	482
HA status.....	483
Session monitoring.....	484
Appendix A: Fortinet MIBs.....	485
Appendix B: Port Numbers.....	487
Appendix C: Scripts.....	489
Events and actions.....	489
Predefined commands.....	489
Control structures.....	496
Operators.....	496
String library.....	498
Examples.....	499
Select content routes based on URI string matches.....	499
Rewrite the HTTP request host header and path.....	500
Rewrite the HTTP response Location header.....	500
Redirect HTTP to HTTPS using Lua string substitution.....	501
Redirect mobile users to the mobile version of a website.....	501
Appendix D: Maximum Configuration Values.....	502
Appendix E: High Speed Logging Binary Format.....	508

Change Log

Date	Change Description
2017-09-21	Initial release.

Introduction

Welcome, and thank you for selecting Fortinet products for your network.

The FortiADC D-series family of application delivery controllers (ADC) optimizes the availability, user experience, performance and scalability of enterprise application delivery.

An ADC is like an advanced server load balancer. An ADC routes traffic to available destination servers based on health checks and load-balancing algorithms; full-featured ADC like FortiADC also improve application performance by assuming some of the server task load. Server tasks that can be handled by the FortiADC appliance include SSL encryption/decryption, WAF protection, Gzip compression, and routing processes, such as NAT.

Features

FortiADC uses Layer 4 and Layer 7 session information to enable an ADC policy and management framework for:

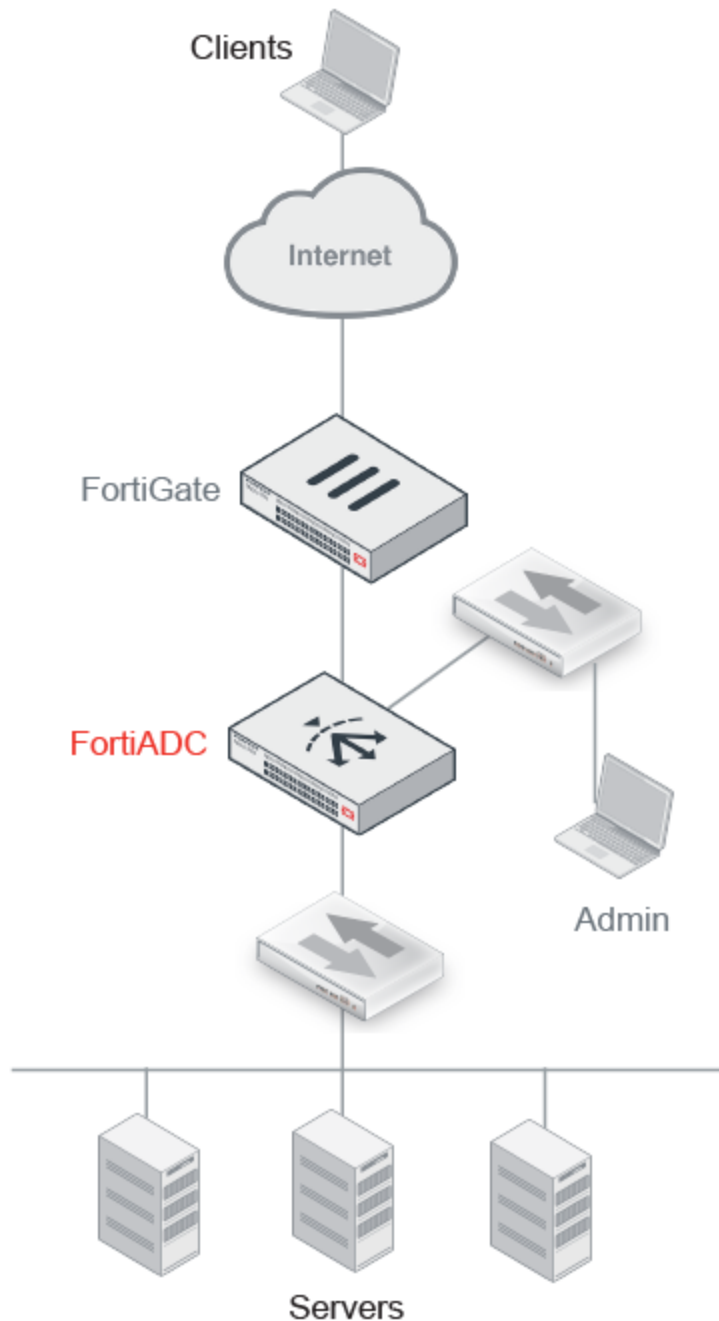
- Server load balancing
- Link load balancing
- Global load balancing
- Security

The FortiADC D-series family includes physical appliances and virtual appliances.

Basic network topology

Your network routing infrastructure should ensure that all network traffic destined for the backend servers is directed to the FortiADC appliance. Usually, clients access backend servers from the Internet through a firewall such as a FortiGate, so the FortiADC appliance should be installed between your servers and the firewall.

[Figure 1](#) shows a basic Router Mode deployment. Refer to the [Basic Deployment Topologies](#) guide for an overview of the packet flow in Router Mode, One-Arm Mode, and Direct Server Return Mode deployments.

Figure 1: Basic network topology

Note: The deployment topology might be different for global load balancing (GLB) or high availability (HA) clusters. Refer to those chapters for a description of features and illustrations.

Scope

This document describes how to use the web user interface to:

- Get started with your deployment.
- Configure feature options.
- Configure network and system settings.
- Monitor the system.
- Troubleshoot issues.

The following topics are covered elsewhere:

- Appliance installation—Refer to the [quick start guide](#) for your appliance model.
- Virtual appliance installation—Refer to the [FortiADC-VM Install Guide](#).
- CLI commands—Refer to the [FortiADC CLI Reference](#). In parts of this manual, brief CLI command examples or CLI syntax are shown to help you understand how the web UI configuration pages are related to the CLI commands.

Chapter 1: What's New

This chapter lists features and enhancements introduced in each of the FortiADC D-Series releases.

FortiADC 4.7.4

- Support for FortiADC 200F hardware model
- Bug fixes

FortiADC 4.7.3

FortiADC 4.7.3 is a patch release only; no new feature or enhancement has been implemented in this release.

FortiADC 4.7.2

FortiADC 4.7.2 offers the following new features or enhancements:

HSM support

- Register HSM server in config file
- Save Client certificate and key to CMDB
- Upload HSM server certificate to FortiADC
- Add registered partition
- Generate CSR with HSM
- View certificate information on the GUI
- Feature configuration supported on both the CLI and the GUI

Support for new hardware models

- FortiADC 1000F
- FortiADC 2000F
- FortiADC 4000F

FortiADC 4.7.1

FortiADC 4.7.1 is a patch release which has fixed some known issues discovered in previous releases. No new features or enhancements have been implemented in this release.

For more information, refer to [FortiADC 4.7.1 Release Notes](#).

FortiADC 4.7.0

Management

- Network Map 2.0
 - Includes SiteMap on link load balance (LLB) and global server load balance (GSLB) modules
- Real server global object
 - Standalone real server objects
 - Allows a single real server to be shared across multiple real server pools and virtual servers
- Configuration templates for Applications
 - Supports SharePoint, Exchange, Windows Remote Desktop, IIS, and Apache

Server load balance (SLB)

- Supports Real-Time Messaging Protocol (RTMP) & Real-Time Streaming Protocol (RTSP)
 - Layer 7 load-balancing
 - Health check
- Supports MySQL
 - Layer 7 load-balancing, user authentication, and persistence
 - Health check
 - MySQL rules
- Decompression
 - Allows decompressed traffic from servers for Layer 7 manipulation (content rewrite), caching, and security (Web Application Firewall)
- Client SSL profile
 - Provides advanced client SSL offloading parameters

User authentication

- Supports LDAP authentication for Regular/Anonymous/LDAPS method
- Supports HTTP basic SSO with HTML Form Authentication/HTML Basic Authentication

High availability (HA)

- Supports HA sync traffic over aggregate ports
- Allows configuration from every device regardless of their HA status (backup vs. master)
- Separated management interface for each node in an HA cluster
- Allows to retrieve license on HA active-passive slave

System

- Transparent mode
 - Support transparent mode installation (Layer 2 forwarding)
- Health check validation

- Allow testing health check policy before bidding it to a real server pool.
- Provide a list of predefined services (TCP, UDP, HTTP, and more)
- Allows to match a admin user to a multiple VDOMs
- Adds Loopback interface in BGB/OSPF defined as router ID
- Attack logs aggregated by date and attack category
- Advanced filters in SLB logs

FortiADC 4.6.2

This is a patch release; no new features or enhancements are implemented. Refer to the [Release Notes](#) for detail.

FortiADC 4.6.1

OpenSSL Library Upgrade

The Software OpenSSL Library has been upgraded to OpenSSL-1.0.2 on FortiADC appliances shipped with the Cavium SSL card, which include the following hardware models:

- FortiADC 400D
- FortiADC 700D
- FortiADC 1500D
- FortiADC 2000D
- FortiADC 4000D

StartTLS

- Supports offloading TLS encryption from back-end SMTP servers

Script

- Supports HTTP:rand_id() function for HTTP

FortiADC 4.6.0

Monitoring and Logs

- Dashboard
- Statistics and information
- Search bar in VS and RS
- Backup server visibility
- Network map
- Three mode views
- Data analytics

DNS load-balancing, security, and caching

- Load-balance DNS traffic (queries and IP addresses) to DNS server
- Sanity check on DNS queries according to RFC 1034, 1035, and 2671
- DNS caching for answer records

Dynamic Load-balancing algorithm

- Dynamic LB based Server Performance such CPU, Memory and Disk

Client certificate forwarding

- Sends client certificates to back-end server for authentication, without affecting SSL offloading

Script validation

- Provides more information in case of syntax error
- Checks content routing for virtual servers
- Generates log message
- Import/export script files

Kerberos Authentication Relay

- Enables authentication between client and server
- Protects against eavesdropping and replay attacks
- Allows nodes communicating over a non-secure network to verify each other's identity in a secure manner

SSL/HTTP visibility (mirroring)

- FortiADC's transparent IP, TCP/S and HTTP/S mirroring capabilities decrypt secure traffic for inspection and reporting by FortiGate or other third-party solutions
- IPv4/IPv6 support

Virtual server port enchantment

- Supports non-consecutive ports in port-range
- Allows Port 0 on TCP or UDP (to catch traffic on all ports)

Security Assertion Markup Language (SAML) 2.0

- Provides Service Provider (SP) and Meta Data of Identity Provider (Idp).
- Can access all VS web resources with user log-in until session expired.

Enhanced Global Load Balancing (GLB) proximity methodology

- Static proximity (GEO, GEO-ISP) and dynamic proximity (RTT, Least Connections, Connection-Limit, Bytes-Per-Second)
- Static match first, dynamic match second

HTTP/S health check

- Adds Username-password Authentication into HTTP/S health check (basic, digest and NTLM)
- Allows to choose SSL Version/Ciphers in HTTPS Health Check

Password policy

- Allows the Admin to control password length and string

VDOM enhancement

- Supports VDOMs restrictions (performance and configuration)
- Able to limit performance (throughput, CPS, SSL, etc.) on each VDOM

SNMP MIBs

- Allows users to download SNMP MIBs from the Web GUI

FortiADC 4.5.3

OpenSSL Library Upgrade

Software OpenSSL library has been upgraded to OpenSSL-1.0.2 on FortiADC appliances shipped with the Cavium SSL card, which include the following hardware models:

- FortiADC 400D
- FortiADC 700D
- FortiADC 1500D
- FortiADC 2000D
- FortiADC 4000D

FortiADC 4.5.2

Software OpenSSL library upgrade

- Software OpenSSL library has been upgraded to openssl-1.0.1s (the latest version) on all FortiADC platforms.
- It's fully functional on FortiADC software.

Enhanced certificate validation

- Support for multiple Online Certificate Status Protocol (OCSP) configurations.
- Support for multiple Certificate Revocation List (CRL) files.

"Description" field for child records in Geo IP Whitelist

- Allows the user to add a brief notation for each child record added to a parent record.

US-Government (USG) mode

- Allows the user to change the appliance from the default regular (REG) mode to USG mode via a special license key.
- Locks the FortiADC D-Series appliance to servers located within the US only.

FortiADC 4.5.1

Acceleration

- Speeds up compression of .PNG, .JPG, and .BMP image files. See
- Caching time definition based on HTTP status code (200/301/302/304)

Server Load Balancing

- SSL Health Check Client certificate selection using SSL Certification
- Support for SIPv6 traffic includes a new health check and virtual server profile
- URL Redirection based on server HTTP status code

High Availability (HA)

- HA-VRRP mode that supports floating IP, traffic group, and fail-over

Global Load Balancing

- Supports DNS SRV record

Miscellaneous

- Full BGP routing support
- Adds a "Description" field in GeoIP White List

FortiADC 4.5.0

SSL offloading

- Support ECDSA SSL cipher suites. See [Chapter 16: SSL Transactions](#).
- SSL certificate validation for server-side SSL connections. See [Configuring real server SSL profiles](#).
- L2 exception list can specify FortiGuard web filter categories. See [Creating a Web Filter Profile configuration](#).

Server Load Balancing

- SIP—Support for SIP traffic includes a new health check, virtual server profile, and persistence method. See [Configuring health checks](#), [Configuring Application profiles](#), and [Configuring persistence rules](#).
- RDP—Support for RDP traffic includes a new virtual server profile and persistence method. See [Configuring Application profiles](#) and [Configuring persistence rules](#).
- HTTP/HTTPS profile—HTTP mode option can be set to HTTP keepalive to support Microsoft SharePoint and other apps that require the session to be kept alive. See [Configuring Application profiles](#).
- Caching—New dynamic caching rules. See [Using caching features](#).
- Real server pool—Member default cookie name is now the real server name. You can change this to whatever you want. See [Using real server pools](#).
- Scripting—Added predefined scripts that you can use as templates. See [Using scripts](#).

Global Load Balancing

- Persistence—Option to enable persistence for specified hosts based on source address affinity. See [Configuring persistence](#).
- Dynamic proximity—Optional configuration for proximity based on least connections. See [Configuring virtual server pools](#).
- Support for @ in zone records. See [Configuring DNS zones](#).
- Zone records (including dynamic records) displayed on zone configuration page. See [Configuring DNS zones](#).

Security

- Bot Detection—Integrated with FortiGuard signatures to allow "good bots" and detect "bad bots." See [Configuring a WAF Profile](#).

Monitoring and Logs

- Fast reports—Real-time statistics and reports for SLB traffic. See [Configuring fast reports](#).
- Session tables and persistence tables—Dashboard tabs for SLB session tables and persistence tables. See [Chapter 20: System Dashboard](#).
- Network map search—Dashboard network map now has search. See [Chapter 20: System Dashboard](#).

System

- New health checks for SIP and custom SNMP. See [Configuring health checks](#).
- Config push/pull (not related to HA). See [Pushing/pulling configurations](#).
- HA sync can be auto/manual. See [Configuring HA settings](#).
- HA status includes details on synchronization. See [Monitoring an HA cluster](#).
- SNMP community host configuration supports subnet address and restriction of hosts to query or trap (or both). See [Configuring SNMP](#).
- Support STARTTLS in email alerts. See [Configuring an SMTP mail server](#).
- Coredump utilities. See [System dump](#).

Platform

- Virtual machine (VM) images for Hyper-V, KVM, Citrix Xen, and opensource Xen. See the [FortiADC-VM Install Guide](#) for details.

FortiADC 4.4.0

Server Load Balancing

- New SSL forward proxy feature can be used to decrypt SSL traffic in segments where you do not have the server certificate and private key. See [Chapter 16: SSL Transactions](#).
- New server-side SSL profiles, which have settings for the FortiADC-to-server connection. This enables you to specify different SSL version and cipher suites for the server-side connection than the ones specified for the client-side connection by the virtual server profile. See [Configuring real server SSL profiles](#).
- Support for ECDHE ciphers, null ciphers, and user-specified cipher lists. See [SSL/TLS versions and cipher suites](#).
- You can now specify a list of SNAT IP address pools in the virtual server configuration. This enables you to use addresses associated with more than one outgoing interface. See [Configuring virtual servers](#).
- Added a health check for UDP, and added hostname to the general settings configuration. In HTTP/HTTPS checks, you can specify hostname instead of destination IP address. See [Configuring health checks](#).
- UDP profiles can now be used with Layer 2 virtual servers. See [Configuring Application profiles](#).
- Server name added to real server pool member configuration. The name can be useful in logs. When you upgrade, the names will be generated from the pool member IP address. You can change that string to whatever you like. See [Using real server pools](#).
- Added a comments setting to the virtual server configuration so you can note the purpose of a configuration. See [Configuring virtual servers](#).

Link Load Balancing

- You can now specify ISP addresses, address groups, and service groups in LLB policies. Using groups adds Boolean OR logic within the elements of LLB rules. See [Configuring link policies](#).

Global Load Balancing

- Added "dynamic proximity" to the server selection algorithm. Dynamic proximity is based on RTT. See [Configuring dynamic proximity](#).
- Added an option to send only a single record in responses instead of an ordered list of records. See [Configuring hosts](#).
- Support for health checks of third-party servers. See [Configuring servers](#).
- Support for TXT resource records. See [Configuring DNS zones](#).

Security

- You can now specify exceptions per WAF profile or per policy. Exceptions identify specific hosts or URL patterns that are not subject to processing by WAF rules. See [Configuring a WAF Profile](#).
- Additional WAF HTTP protocol constraint rules. See [Configuring a WAF Profile](#).

Monitoring and Logs

- Added a Network Map tab to the dashboard. In the Network Map, each virtual server is a tree. The status of the virtual server and real server pool members is displayed. See [Chapter 20: System Dashboard](#).
- Added on-demand and scheduled reports for many common queries. You can also configure custom queries. See [Configuring reports](#).
- Added event log categories and added a column in logs to support future integration with FortiAnalyzer. Removed the Download Logs page. Each log category page now has a **Download** button. See [Using the event log](#).
- Enhanced SNMP MIBs and traps. See [Appendix A: Fortinet MIBs](#) for information on downloading the vendor-specific and product-specific MIB files.

System

- Shared Resources—Merged the address and service configuration for firewall and LLB. Added address groups and service groups, which can be used in LLB policy rules. See [Chapter 10: Shared Resources](#).
- Routing—Support for OSPF authentication. See [OSPF](#).
- HA—Added option to actively monitor remote beacon IP addresses to determine if the network path is available. See [Configuring HA settings](#).
- System—Updated the web UI to match CLI configuration options for global administrator and access profile. See [Managing administrator users](#).
- Web UI—Support for Simplified Chinese. See [Configuring basic system settings](#).
- Troubleshooting—New commands: `diagnose debug flow`, `diagnose debug report`, `diagnose debug timestamp`, `execute checklogdisk`, and `execute fixlogdisk`. See the [FortiADC CLI Reference](#).
- CLI—Added `execute ssh` and `execute telnet` for connections to remote hosts.

API

- REST API—Remote configuration management with a REST API. See the [FortiADC REST API Reference](#).

FortiADC 4.3.1

- Server Load Balancing Persistence—Added a Match Across Servers option to the Source Address affinity method. This option is useful when the client session for an application has connections over multiple ports (and thus multiple virtual servers). This option ensures the client continues to access the same backend server through different virtual servers for the duration of a session.

- Server Load Balancing TCP Multiplexing— Added support for HTTPS connections.
- Global Load Balancing DNS Server—The negative caching TTL in the SOA resource record is now configurable.

FortiADC 4.3.1

- Virtual domains—Increased the maximum number of VDOMs on the following platforms:
 - FortiADC 700D — 30
 - FortiADC 1500D — 45
 - FortiADC 2000D — 60
 - FortiADC 4000D — 90
- Health checks—Added an HTTP Connect health check that is useful for testing the availability of web cache proxies, such as FortiCache.
- ISP address book—Added a province location setting to the ISP address book. The province setting is used in GLB deployments in China to enable location awareness that is province-specific. For example, based on location, the DNS server can direct a user to a datacenter in Beijing or Guangdong rather than the broader location China. Only a predefined set of Chinese provinces is supported.
- Advanced routing—Exception list for reverse path route caching.

FortiADC 4.3.0

- Authentication—Framework to offload authentication from backend servers.
- Geo IP blocking—Policy that takes the action you specify when the virtual server receives requests from IP addresses in the blocked country's IP address space.
- Web application firewall—Protect against application layer attacks with policies such as signatures, HTTP protocol constraints, request URL and file extension patterns, and SQL/XSS injection detection.
- Scripts—Support for Lua scripts to perform actions that are not currently supported by the built-in feature set.
- SSL/TLS—Support for PFS ciphers.
- Health check improvements—The SLB and LLB health check configuration has been combined and moved to System > Shared Resources. You can configure destination IP addresses for health checks. This enables you to test both the destination server and any related services that must be up for the server to be deemed available. Also added support for Layer 2 and SSH health checks.
- Port range—Support for virtual IP address with a large number of virtual ports.
- NAT46/64—Support for NAT46/64 by the SLB module.
- ISP address book—Framework for an ISP address book that simplifies the ISP route and LLB proximity route configuration.
- Proximity routes—Support for using ISP address book entries in the LLB proximity route table.
- Backup pool member—Support for designating a link group or virtual tunnel group member as a “backup” that joins the pool when all of the main members are unavailable.
- Global load balancing—New framework that leverages the FortiGuard Geolocation database or the FortiADC predefined ISP address books to direct clients to the closest available FortiADC virtual servers.
- Stateful firewall—If client-to-server traffic is allowed, the session is maintained in a state table, and the response traffic is allowed.
- Virtual server traffic—Many of the firewall module features can be applied to virtual server traffic.

- **ISP Routes**—ISP routes are used for outbound traffic and link load balancing traffic.
- **HA upgrade**—Simpler one-to-many upgrade from the primary node.
- **HA status**—HA status tab on the system dashboard.
- **HA remote login**—You can use the `execute ha manage` command to connect to the command-line interface of a member node. See the CLI reference.
- **SNMPv3 support**
- **Statistics and log database** to better support dashboard and report queries.
- **Improved dashboard**—New time period options for the virtual server throughput graphs.
- **Improved reports**—New report queries for SLB HTTP virtual server reports, including client IP address, client browser type, client OS, and destination URL.
- **Backup & restore**—Option to back up the entire configuration, including error page files, script files, and ISP address books.

New CLI commands to facilitate troubleshooting:


- `diagnose debug config-error-log`—Use this command to see debug errors that might be generated after an upgrade or major configuration change.
- `diagnose debug crashlog`—Use this command to manage crashlog files. Typically, you use these commands to gather information for Fortinet Services & Support.
- `execute statistics-db`—Use this command to reset or restore traffic statistics.
- `config system setting`—Use this command to configure log database behavior (overwrite or stop writing) when disk utilization reaches its capacity.

For details, see the CLI reference.

FortiADC 4.2.3

- **HTTPS and TCPS Profiles**—Support for SHA-256 ciphers suites.

FortiADC 4.2.2

- **Content rewriting**—Support for PCRE capture and back reference to write the Location URL in redirect rules.
- **Web UI**—You can clone configuration objects to quickly create similar configuration objects. If a configuration object can be cloned, the copy icon  appears in the tools column for its summary configuration page.
- **Web UI**—You can sort many of the configuration summary tables by column values. If a configuration summary table can be sorted, it includes sort arrows in the column headings. For example, the Server Load Balance > Virtual Server configuration summary page can be sorted by Availability, Status, Real Server pool, and so on. You can also sort the Dashboard > Virtual Server > Real Server list by column values—for example, by Availability, Status, Total Sessions, or throughput bytes.

FortiADC 4.2.1

Bug fixes only.

FortiADC 4.2.0

- New web UI
- New log subtypes
- New dashboard and report features
- Additional load balancing methods—Support for new methods based on a hash of a full URI, domain name, hostname, or destination IP address.
- Predefined health checks—Helps you get started with your deployment.
- Predefined persistence rules—Helps you get started with your deployment.
- HTTP Turbo profile—Improves the performance of HTTP applications that do not require our optional profile features.
- Layer 2 load balancing—Support for TCP profiles.
- Granular SSL configuration—Specify the SSL/TLS versions and encryption algorithms per profile.
- Connection rate limiting—Set a connection rate limit per real server or per virtual server.
- HTTP transaction rate limiting—Set a rate limit on HTTP transactions per virtual server.
- Additional link load balancing methods—Support for new methods in link groups, including spillover and hash of the source IP address.
- Global load balancing—A new implementation of our DNS-based solution that enables you to deploy redundant resources around the globe that you can leverage to keep your business online when a local area deployment experiences unexpected spikes or downtime.
- HA active-active clustering—Support for active-active clusters.
- Administrator authentication enhancements—Support for authenticating users against LDAP and RADIUS servers.
- Multinetting—You can configure a secondary IP address for a network interface when necessary to support deployments with backend servers that belong to different subnets.
- High speed logging—Supports deployments that require a high volume of logging activity.
- Packet Capture—Support for tcpdump.

FortiADC 4.1

No design changes. Bug fixes only.

FortiADC 4.0 Patch 2

No design changes. Bug fixes only.

FortiADC 4.0 Patch 1

No design changes. Bug fixes only.

FortiADC 4.0

- VDOMs—Virtual domains (VDOMs) allow you to divide a FortiADC into two or more virtual units that are configured and function independently. The administrator for each virtual domain can view and manage the configuration for his or her domain. The `admin` administrator has access to all virtual domain configurations.
- Caching – A RAM cache is a cache of HTTP objects stored in FortiADC's system RAM that are reused by subsequent HTTP transactions to reduce the amount of load on the backend servers.
- IP Reputation—You can now block source IP addresses that have a poor reputation using data from the FortiGuard IP Reputation Service.
- Layer 2 server load balancing—FortiADC can now load balance Layer 3 routers, gateways or firewalls. This feature is useful when the request's destination IP is unknown and you need to load balance connections between multiple next-hop gateways. Supports HTTP, HTTPS and TCPS client-side connection profiles only.
- Open Shortest Path First (OSPF) support—The new OSPF feature allows FortiADC to learn dynamic routes from or redistribute routes to neighboring routers.
- HTTPS profile type for virtual servers—The HTTPS profile type provides a standalone HTTPS client-side connection profile.
- Consistent Hash IP – The persistence policy type Hash IP has changed to Consistent Hash IP. Consistent hashing allows FortiADC to achieve session persistence more efficiently than traditional hashing.
- Enhanced logs

FortiADC 3.2.0

- Link routing policies—You can now specify how FortiADC routes traffic for each available ISP link, including by source or destination address and port.
- Virtual tunnels—You can now use tunneling between two FortiADC appliances to balance traffic across multiple links to each appliance. A typical scenario is a VPN between a branch office and headquarters for application-specific access.
- Persistent routing—You can now configure connections that persist regardless of the FortiADC link load balancing activity. You can configure persistence based on source IP, destination IP, and subnet.
- Proximity-based routing—Maximize WAN efficiency by using link proximity to determine latency between FortiADC and remote WAN sites so that FortiADC can choose the best route for traffic.
- Scheduled link load balancing—You can now apply a link load balancing policy during a specific time period.
- One-to-one (1-to-1) NAT—You can now fully define how each individual source and destination IP address will be translated. This feature is useful when you require a different NAT range for each ISP.
- PPPoE interface support—To support DSL connectivity, you can now configure interfaces to use PPPoE (Point-to-Point Protocol over Ethernet) to automatically retrieve its IP address configuration.

FortiADC 3.1.0

- Custom error page—You can now upload a custom error page to FortiADC that it can use to respond to clients when HTTP service is unavailable.

- Full NAT for Layer 3/4 load balancing—Layer 3/4 load balancing now supports full NAT (translation of both source and destination IP addresses). FortiADC can now round robin among a pool of source IP addresses for its connections to backend servers.
- Standby server—You can now configure FortiADC to forward traffic to a hot standby (called a Backup Server) when all other servers in the pool are unavailable.
- Log cache memory—To avoid hard disk wear and tear, FortiADC can cache logs in memory and then periodically write them to disk in bulk. Previously, FortiADC always wrote each log message to disk instantaneously.
- HA sync for health check status with IPv6—For high availability FortiADC clusters, the Layer 4 health check status of IPv6-enabled virtual servers is now synchronized.

FortiADC 3.0.0

- Link load balancing—FortiADC now supports load balancing among its links, in addition to distributing among local and globally distributed servers. Depending on if the traffic is inbound or outbound, different mechanisms are available: outbound can use weighted round robin; inbound can use DNS-based round robin or weighted round robin.
- HTTP response compression—FortiADC now can compress responses from your backend servers, allowing you to off load compression from your backend servers for performance tuning that delivers faster replies to clients.
- Quality of service (QoS)—FortiADC now can guarantee bandwidth and queue based upon source/destination address, direction, and network service.
- Source NAT (SNAT)—When applying NAT, FortiADC can now apply either static or dynamic source NAT, depending on your preference.
- Session persistence by source IP segment—FortiADC now can apply session persistence for entire segments of source IPs such as 10.0.2.0/24. Previously, session persistence applied to a single source IP.
- Health check enhancements—FortiADC now supports additional health check types for servers that respond to these protocols: email (SMTP, POP3, IMAP), TCPS, TCP `SYN` (half-open connection), SNMP, and UDP.
- HA enhancements—FortiADC HA now synchronizes Layer 3/4 and Layer 7 sessions and connections for session persistence and uninterrupted connections when the standby assumes control of traffic.

FortiADC 2.1.0

Support for FortiADC 200D and FortiADC VM—FortiADC software has been released to support these new platforms.

Chapter 2: Key Concepts and Features

This chapter includes the following topics:

- [Server load balancing](#)
- [Link load balancing](#)
- [Global load balancing](#)
- [Security](#)
- [High availability](#)
- [Virtual domains](#)

Server load balancing

Server load balancing (SLB) features are designed to give you flexible options for maximizing performance of your backend servers. The following topics give an overview of SLB features:

- [Feature Summary](#)
- [Authentication](#)
- [Caching](#)
- [Compression](#)
- [Content rewriting](#)
- [Content routing](#)
- [Scripting](#)
- [SSL transactions](#)

Feature Summary

Table 1 summarizes server load balancing features.

Table 1: Server load balancing features

Features	Summary
Methods	<ul style="list-style-type: none">• Round robin• Weighted round robin• Least connections• Fastest response• Hash of URI, domain, host, destination IP
Health check	Checks based on Layer 3, Layer 4, or Layer 7 data.

Features	Summary
Server management	<ul style="list-style-type: none"> • Warm up • Rate limiting • Maintenance mode with session ramp down
Persistence	Based on: <ul style="list-style-type: none"> • Cookies • TCP/IP header matches • A hash of TCP/IP header values • TLS/SSL session ID • RADIUS attribute • RDP Session Broker cookie • SIP caller ID
Layer 7	Profiles: HTTP, HTTPS, HTTP Turbo, RADIUS, RDP, SIP, TCPS Content routing: HTTP Host, HTTP Referer, HTTP Request URL, SNI hostname, Source IP address Content rewriting: URL redirect, 403 Forbidden, or HTTP request/response rewrite
Layer 4	Profiles: FTP, TCP, UDP Content routing: Source IP address
Layer 2	Profiles: HTTP, HTTPS, TCP, TCPS, UDP Note: Layer 2 load balancing is useful when the request's destination IP is unknown and you need to load balance connections between multiple next-hop gateways.

For detailed information, see [Chapter 4: Server Load Balancing](#).

Authentication

FortiADC SLB supports offloading authentication from backend servers. The auth policy framework supports authentication against local, LDAP, and RADIUS authentication servers, and it enables you to assign users to groups that are authorized to access protected sites.

For configuration details, see [Configuring authentication policies](#).

Caching

FortiADC SLB supports both static and dynamic caching. Caching reduces server overload, bandwidth saturation, high latency, and network performance issues.

When caching is enabled for a virtual server profile, the FortiADC appliance dynamically stores application content such as images, videos, HTML files and other file types to alleviate server resources and accelerate overall application performance.

For configuration details, see [Using caching features](#).

Compression

FortiADC SLB supports compression offloading. Compression offloading means the ADC handles compression processing instead of the backend servers, allowing them to dedicate resources to their own application processes.

When compression is enabled for a virtual server profile, the FortiADC system intelligently compresses HTTP and HTTPS traffic. Reducing server reply content size accelerates performance and improves response times. FortiADC supports both industry standard GZIP and DEFLATE algorithms.

For configuration details, see [Configuring compression rules](#).

Decompression

FortiADC SLB also supports decompression of HTTP request body before sending it to the Web Application Firewall (WAF) for scanning according to the content-encoding header. Upon receiving a compressed HTTP request body, FortiADC first uses the zlib library to extract the HTTP body to a temporary buffer and then sends the buffer to the WAF engine for scanning.

Content rewriting

FortiADC SLB supports content rewriting rules that enable you to rewrite HTTP requests and responses so that you can cloak the details of your internal network. You can also create rules to redirect requests.

For configuration details and examples, see [Using content rewriting rules](#).

Content routing

FortiADC SLB supports content routing rules that direct traffic to backend servers based on source IP address or HTTP request headers.

For configuration details, see [Configuring content routes](#).

Scripting

FortiADC SLB supports Lua scripts to perform actions that are not currently supported by the built-in feature set. Scripts enable you to use predefined script commands and variables to manipulate the HTTP request/response or select a content route.

For configuration details, see [Using scripts](#).

SSL transactions

FortiADC SLB supports SSL offloading. SSL offloading means the ADC handles SSL decryption and encryption processing instead of the backend servers, allowing the backend servers to dedicate resources to their own application processes.

SSL offloading results in improved SSL/TLS performance. On VM models, acceleration is due to offloading the cryptographic processes from the backend server. On hardware models with ASIC chips, cryptography is also hardware-accelerated: the system can encrypt and decrypt packets at better speeds than a backend server with a general-purpose CPU.

FortiADC SLB also supports SSL decryption by forward proxy in cases where you cannot copy the server certificate and private key to the FortiADC, either because it is impractical or impossible (in the case of outbound traffic to unknown Internet servers).

For detailed information, see [Chapter 16: SSL Transactions](#).

Link load balancing

Link load balancing (LLB) features are designed to manage traffic over multiple ISP or WAN links. This enables you to provision multiple links, resulting in reduced risk of outages and additional bandwidth to relieve traffic congestion.

For detailed information, see [Chapter 5: Link Load Balancing](#).

Global load balancing

Global load balancing (GLB) makes your network reliable and available by scaling applications across multiple data centers to improve application response times and be prepared for disaster recovery.

You can deploy DNS to direct traffic based on application availability and location.

For detailed information, see [Chapter 6: Global Load Balancing](#).

Security

In most deployment scenarios, we recommend you deploy FortiGate to secure your network. Fortinet includes security functionality in the FortiADC system to support those cases when deploying FortiGate is impractical. FortiADC includes the following security features:

- Firewall—Drop traffic that matches a source/destination/service tuple you specify.
- Security connection limit—Drop an abnormally high volume of traffic from a source/destination/service match.
- IP Reputation service—Drop or redirect traffic from source IPs that are on the FortiGuard IP Reputation list.
- Geo IP—Drop or redirect traffic from source IPs that correspond with countries in the FortiGuard Geo IP database.
- Web application firewall—Drop or alert when traffic matches web application firewall attack signatures and heuristics.
- Denial of service protection—Drop half-open connections to protect the system from a SYN flood attack.

For detailed information, see [Chapter 7: Network Security](#).

High availability

The FortiADC appliance supports high availability features like active-passive, active-active cluster, active-active-VRRP cluster, failure detection, and configuration synchronization. High availability deployments can support 99.999% service level agreement uptimes. For detailed information, see [Chapter 14: High Availability Deployments](#).

Virtual domains

A virtual domain (VDM) is a complete FortiADC instance that runs on the FortiADC platform. The VDM feature supports multitenant deployments. To do this, you create a virtual domain configuration object that contains all of the system and feature configuration options of a full FortiADC instance, and you provision an administrator account with privileges to access and manage only that VDM. For detailed information, see [Chapter 15: Virtual Domains](#).

Chapter 3: Getting Started

This chapter provides the basic workflow for getting started with a new deployment.

Basic steps:

1. Install the appliance.
2. Configure the management interface.
3. Configure the following basic network settings:
 - Administrator password
 - System date and time
 - Network interfaces
 - DNS
4. Test connectivity.
5. Complete product registration, install your license, and update the firmware.
6. Configure a basic load balancing policy.
7. Test the deployment with load to verify expected behavior.
8. Back up this basic configuration so that you have a restore point.



Tips:

- Configuration changes are applied to the running configuration as soon as you save them.
 - Configuration objects are saved in a configuration management database. You cannot change the name of a configuration object after you have initially saved it.
 - You cannot delete a configuration object that is referenced in another configuration object (for example, you cannot delete an address if it is used in a policy).
-

Step 1: Install the appliance

This Handbook assumes you have already installed the appliance into a hardware rack or the virtual appliance into a VMware environment.

For information on hardware appliances, refer to the FortiADC hardware manuals.

For information on the virtual appliance, refer to the FortiADC-VM Install Guide.

To download these documents, go to:

<http://docs.fortinet.com/fortiadc-d-series/hardware>

Step 2: Configure the management interface

You use the management port for administrator access. It is also used for management traffic (such as SNMP or syslog). If your appliance has a dedicated management port, that is the port you configure as the management interface; otherwise, it is the convention to use port1 for the management interface.

You configure the following basic settings to get started so that you can access the web UI from a remote location (like your desk):

- **Static route**—Specify the gateway router for the management subnet so you can access the web UI from a host on your subnet.
- **IP address**—You typically assign a static IP address for the management interface. The IP address is the host portion of the web UI URL. For example, the default IP address for the management interface is 192.168.1.99 and the default URL for the web UI is <https://192.168.1.99>.
- **Access**—Services for administrative access. We recommend HTTPS, SSH, SNMP, PING.

Before you begin:

- You must know the IP address for the default gateway of the management subnet and the IP address that you plan to assign the management interface.
- You need access to the machine room in which a physical appliance has been installed. With physical appliances, you must connect a cable to the management port to get started.
- You need a laptop with an RJ-45 Ethernet network port, a crossover Ethernet cable, and a web browser (a recent version of Chrome, Firefox, or Internet Explorer).
- Configure the laptop Ethernet port with the static IP address 192.168.1.2 and a netmask of 255.255.255.0. These settings enable you to access the FortiADC web UI as if from the same subnet as the FortiADC in its factory configuration state.

To connect to the web UI:

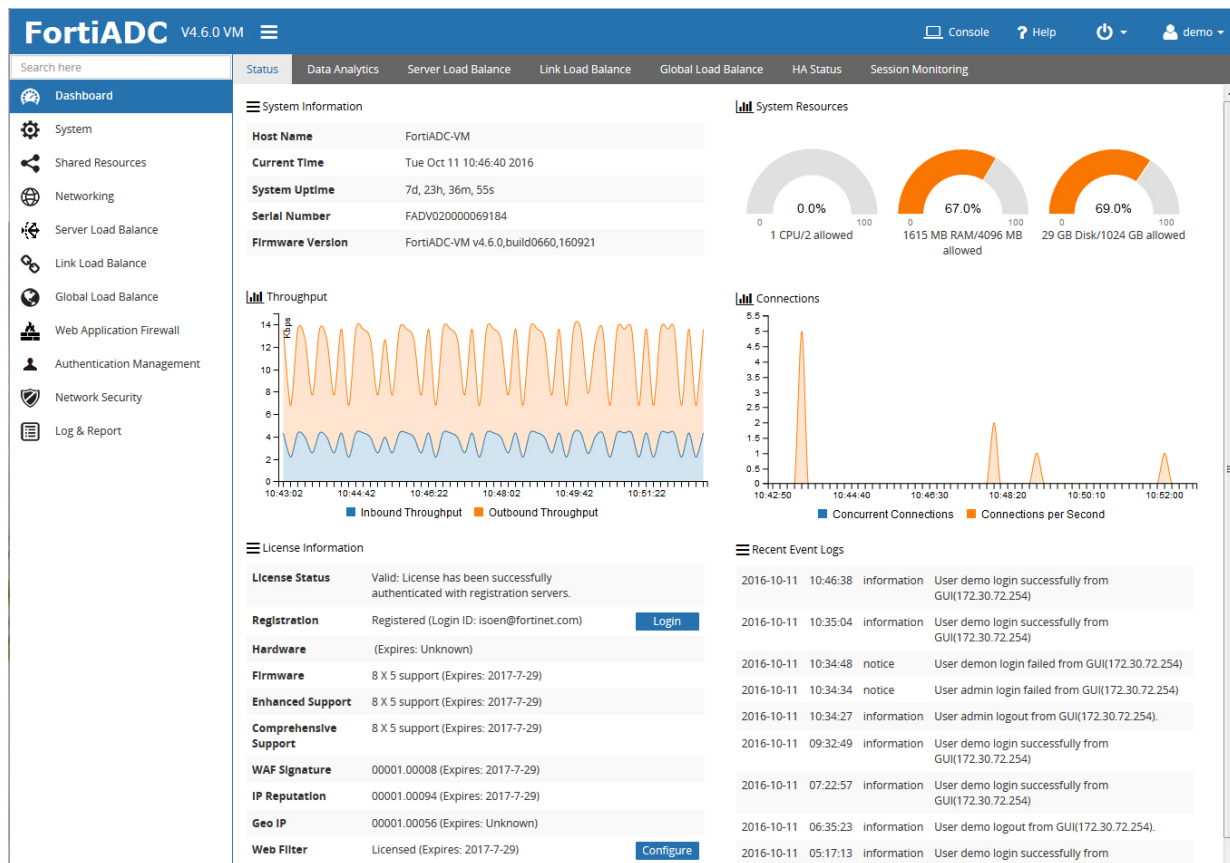
1. Use the crossover cable to connect the laptop Ethernet port to the FortiADC management port.
2. On your laptop, open the following URL in your web browser:
<https://192.168.1.99/>

The system presents a self-signed security certificate, which it presents to clients whenever they initiate an HTTPS connection to it.

3. Verify and accept the certificate, and acknowledge any warnings about self-signed certificates.
The system displays the administrator login page. See [Figure 2](#).

Figure 2: Login pageThe image shows the login page for FortiADC V4.6.0 VM. It features a blue header with the text "FortiADC" on the left and "V4.6.0 VM" on the right. Below the header, there are two input fields: "Username" and "Password". Below these fields is a large blue button with the text "Log In" in white.

4. Enter the username **admin** and no password.
The system displays the dashboard. See [Figure 3](#).

Figure 3: Dashboard after initial login

To complete the procedures in this section using the CLI:

1. Use an SSH client such as PuTTY to make an SSH connection to 192.168.1.99 (port 22).
2. Acknowledge any warnings and verify and accept the FortiADC SSH key.
3. Enter the username **admin** and no password.
4. Use the following command sequence to configure the static route:

```
config router static
  edit 1
    set gateway <gateway_ipv4>
  end
end
```

5. Use the following command sequence to configure the management interface:

```
config system interface
  edit <interface_name>
    set ip <ip&netmask>
    set allowaccess {http https ping snmp
      ssh telnet}
  end
end
```

The system processes the update and disconnects your SSH session because the interface has a new IP address. At this point, you should be able to connect to the CLI from a host on the management subnet you just configured. You can verify the configuration remotely.



Step 3: Configure basic network settings

The system supports network settings for various environments.


To get started, you configure the following basic settings:

- Administrator password—You must change the password for the **admin** account.
- System date and time—We recommend you use NTP to maintain the system time.
- Network interfaces—You must configure interfaces to receive and forward the network traffic to and from the destination servers.
- DNS—You must specify a primary and secondary server for system DNS lookups.

Before you begin:

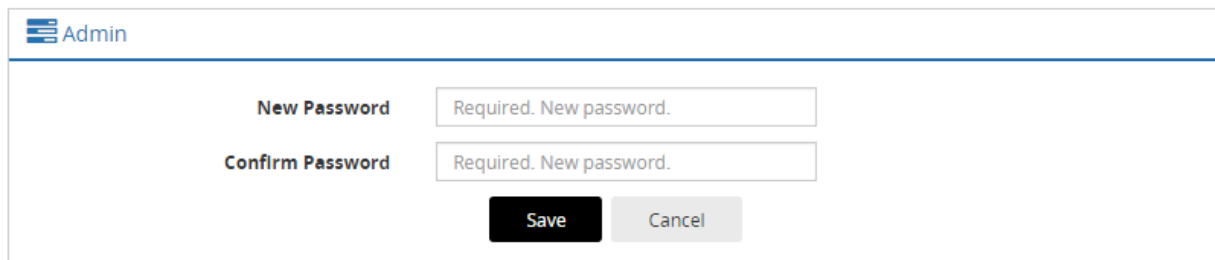
- You must know the IP address for the NTP servers your network uses to maintain system time.
- You must know the IP addresses that have been provisioned for the traffic interfaces for your FortiADC deployment.
- You must know the IP address for the primary and secondary DNS servers your network uses for DNS resolution.

To change the admin password:

1. Go to System > Administrator to display the configuration page.
2. Double-click the key icon  in the row for the user **admin** to display the change password editor. See [Figure 4](#).
3. Change the password and save the configuration.

For detailed information on configuring administrator accounts, refer to the online help or see [Managing administrator users](#).

Figure 4: System administrator change password editor



CLI commands:



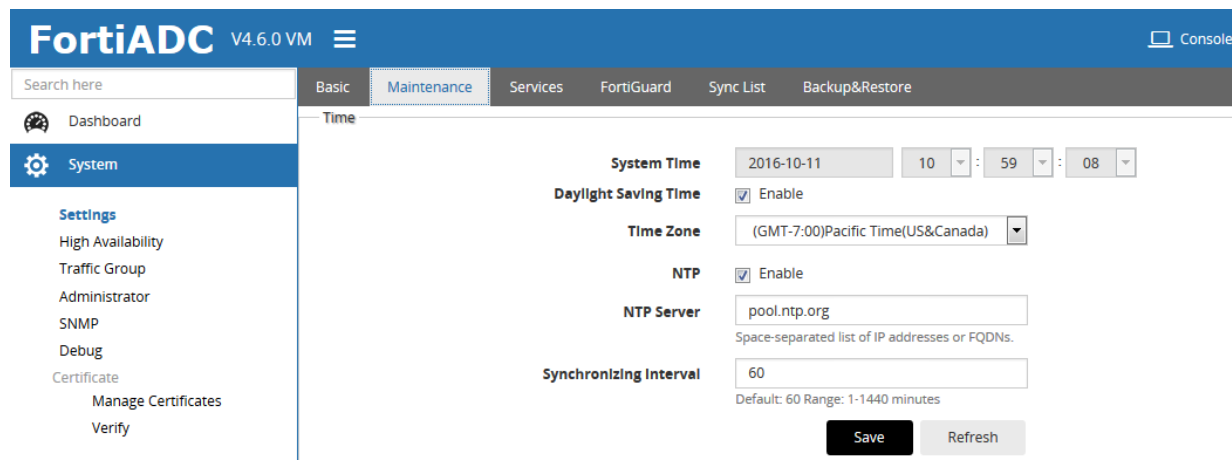
```
FortiADC-VM # config system admin
FortiADC-VM (admin) # edit admin
FortiADC-VM (admin) # set password <string>
Current password for 'admin':
FortiADC-VM (admin) # end
```

To configure system time:

1. Go to System > Settings.
2. Click the **Maintenance** tab to display the configuration page. See [Figure 5](#).
3. Enter NTP settings and save the configuration.

For detailed information, refer to the online help or see [Configuring system time](#).

Figure 5: System time configuration page



**CLI commands:**

```
config system time ntp
set ntpsync enable
set ntpserver {<server_fqdn> | <server_ipv4>}
set syncinterval <minutes_int>
end
```

Or use a command syntax similar to the following to set the system time manually:

```
config system time manual
set zone <timezone_index>
set daylight-saving-time {enable | disable}
end
execute date <MM/DD/YY> <HH:MM:SS>
```

To configure network interfaces:

1. Go to Networking > Interface to display the configuration page.
2. Double-click the row for port2, for example, to display the configuration editor. See [Figure 6](#).
3. Enter the IP address and other interface settings and save the configuration.

For detailed information, refer to the online help or see [Configuring network interfaces](#).

Figure 6: Network interface configuration page

Interface

Name: port2

Status: ☒ Up ☐ Down

Allow Access: ☒ HTTPS ☒ Ping ☒ SSH ☒ SNMP ☒ HTTP ☒ Telnet

Type: Physical

Mode: ☒ Static ☐ PPPoE

Traffic Group: default

Floating: ☐ Enable

Mode Specifics

IPv4/Netmask: 10.1.50.10/24
Example: 192.0.2.5/24

IPv6/Netmask: ::/0
Example: 2001:0db8:85a3::8a2e:0370:7334/64

Secondary IP Address: ☐ Enable

HA Node IP Address List

ID	IP Address	Node ID	Allow Access
Total: 0 Refresh			

Save Cancel

CLI commands:



```

config system interface
  edit <interface_name>
    set ip <ip&netmask>
    set allowaccess {http https ping snmp ssh
telnet}
  end
end

```

To configure DNS:

1. Go to System > Settings to display the Basic configuration page. See [Figure 7](#).
2. Enter the IP address for a primary and secondary DNS server; then save the configuration.

For detailed information on configuring DNS, refer to the online help or see [Configuring basic system settings](#).

Figure 7: DNS configuration page

The screenshot shows the FortiADC web interface for DNS configuration. The left sidebar contains navigation links: Dashboard, System (selected), Settings (with sub-links: High Availability, Traffic Group, Administrator, SNMP, Debug, Certificate, Manage Certificates, Verify), Shared Resources, and Networking. The main content area has tabs: Basic, Maintenance, Services, FortiGuard, Sync List, and Backup&Restore. The Basic tab is active, showing the following configuration fields:

- Hostname:** FortiADC-VM
- Language:** English (dropdown menu)
- Idle Timeout:** 30 (Default: 30 Range: 1-480 minutes)
- HTTP Port:** 80 (Default: 80 Range: 1-65535)
- HTTPS Port:** 443 (Default: 443 Range: 1-65535)
- SSH Port:** 22 (Default: 22 Range: 1-65535)
- Telnet Port:** 23 (Default: 23 Range: 1-65535)
- Primary DNS:** 208.91.112.53
- Secondary DNS:** 208.91.112.52
- Virtual Domain:** ☐ Enable
- Config Sync:** ☐ Enable

At the bottom right, there are 'Save' and 'Refresh' buttons.



CLI commands:

```
config system dns
  set primary <address_ipv4>
  set secondary <address_ipv4>
end
```

Step 4: Test connectivity to destination servers

Use ping and traceroute to test connectivity to destination servers.

To test connectivity from the FortiADC system to the destination server:

Run the following commands from the CLI:

```
execute ping <destination_ip4>
execute traceroute <destination_ip4>
```

To test connectivity from the destination server to the FortiADC system:

1. Enable ping on the network interface.
2. Use the ping and traceroute utilities available on the destination server to test connectivity to the FortiADC network interface IP address.

For troubleshooting tips, see [Chapter 19: Troubleshooting](#).

Step 5: Complete product registration, licensing, and upgrades

Your new FortiADC appliance comes with a factory image of the operating system (firmware). However, if a new version has been released since factory imaging, you might want to install the newer firmware before continuing the system configuration.

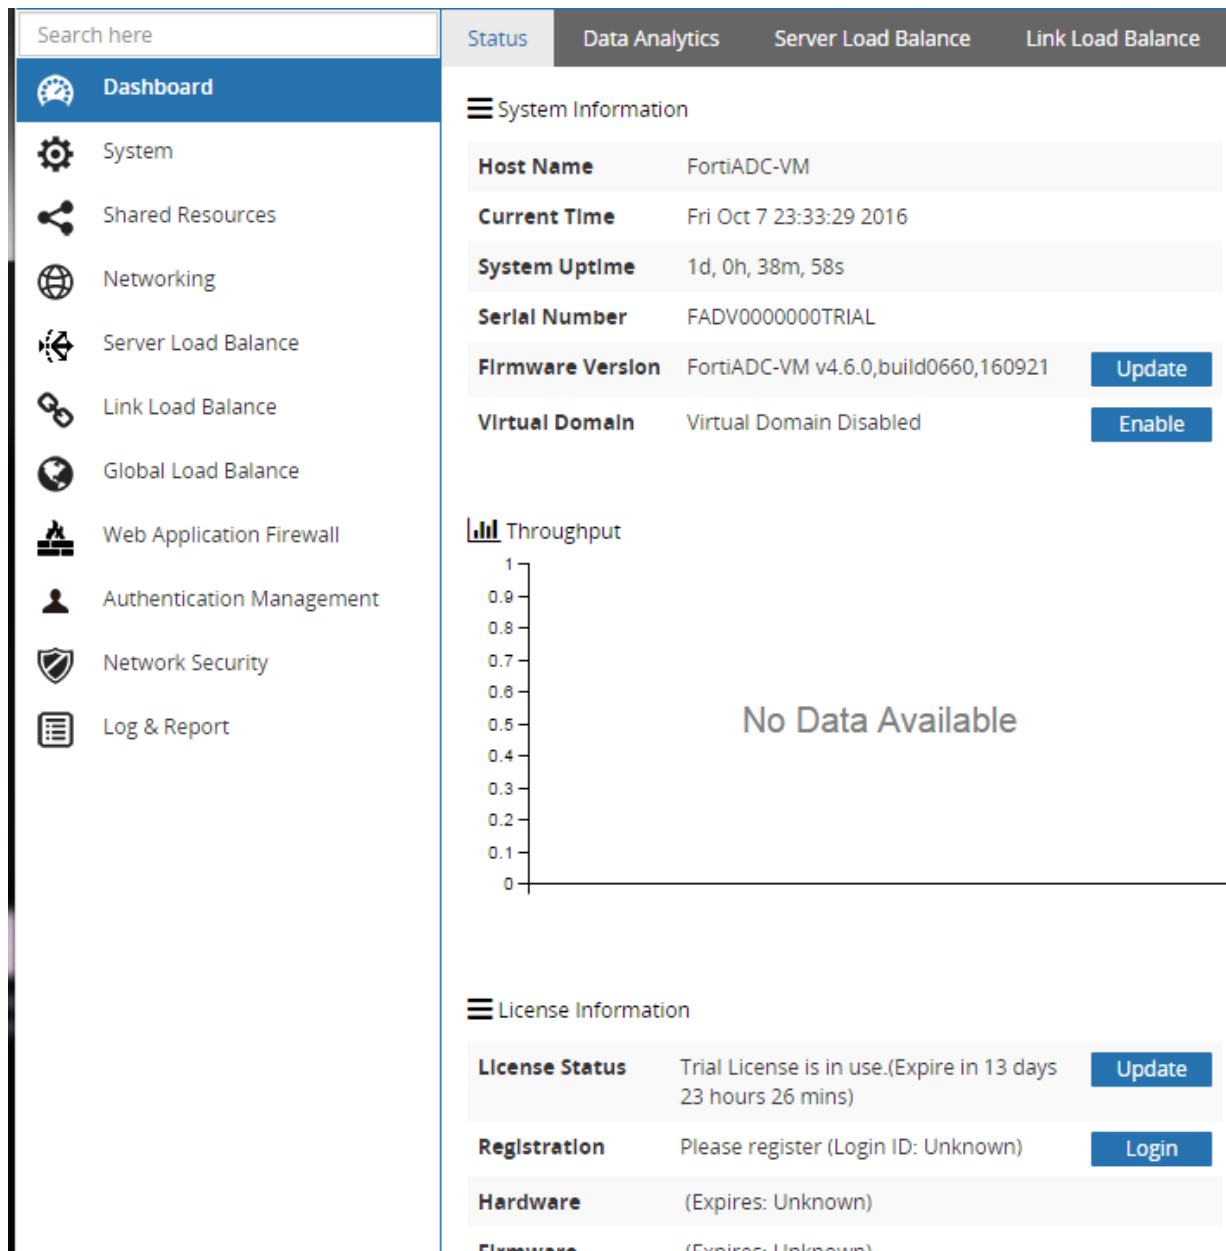
Before you begin:

- Register—Registration is required to log into the Fortinet Customer Service & Support site and download firmware upgrade files. For details, go to <http://kb.fortinet.com/kb/documentLink.do?externalID=12071>.
- Check the installed firmware version—Go to the dashboard. See [Figure 8](#).
- Check for upgrades—Major releases include new features, enhancements, and bug fixes. Patch releases can include enhancements and bug fixes.
- Download the release notes at <http://docs.fortinet.com/fortiadc-d-series/>.
- Download firmware upgrades at <https://support.fortinet.com/>.

To upload your license and new firmware:

1. Go to the dashboard.
2. Under License Status, click **Update** to locate and upload the license file.
3. Under Firmware Version, click **Update** to locate and upload the firmware file.

For detailed information, refer to the online help or see [Updating firmware](#).

Figure 8: License and firmware upgrade page

Step 6: Configure a basic server load balancing policy

A FortiADC server load balancing policy has many custom configuration options. You can leverage the predefined health check, server profile, and load balancing method configurations to get started in two basic steps:

1. Configure the real server pool.
2. Configure the virtual server features and options.

For complete information on server load balancing features, start with [Server load balancing basics](#).

To configure the server pool:

1. Go to Server Load Balance > Real Server to display the configuration page.
2. Click **Add** to display the configuration editor. See [Figure 9](#).
3. Complete the basic configuration and click **Save**.
4. Double-click the configuration to display the configuration editor.
5. Under Member, click **Add** to display the Edit Member configuration editor. See [Figure 10](#).
6. Complete the member configuration and click **Save**.

For detailed information, refer to the online help or see [Configuring real server pools](#).

Figure 9: Real server pool basic configuration page

Real Server Pool

Name Required config name. No spaces.

Address Type ☒ IPv4 ☐ IPv6

Health Check ☐ Enable

Real Server SSL Profile NONE

Member

+ Add X Delete Total: 0 Refresh

Please save parent record first!

ID	Name	Address	Health Check	Port	Weight	Limit	Status	Backup	

Save Cancel

Figure 10: Real server pool member configuration page

Real Server Pool

Real Server Pool **Edit Member**

Status ☒ Enable ☐ Disable ☐ Maintain

Server Name Required. Specify the servi

Address 0.0.0.0
Example: 192.0.2.1

Port 80
Default: 80 Range: 0-65535

Weight 1
Default: 1 Range: 1-256

Recover 0
Default: 0 (disabled) Range: 0-86400 seconds

Warm Up 0
Default: 0 (disabled) Range: 0-86400 seconds

Warm Rate 100
Default: 10 Range: 1-86400 connections per second

Connection Limit 0
Default: 0 (disabled) Range: 0-1048576 concurrent connections

Cookie Please input cookie

Connection Rate Limit 0
Default: 0 (disabled) Range: 0-86400 connections per second

Health Check Inherit ☒ Enable

Backup ☐ Enable

RS Profile Inherit ☒ Enable

Save **Cancel**

To configure the virtual server:

1. Go to Server Load Balance > Virtual Server to display the configuration page.
2. Click **Add** to display the configuration editor. See [Figure 11](#).
3. Complete the configuration and click **Save**.

For detailed information, refer to the online help or see [Configuring virtual servers](#).

Figure 11: Virtual server configuration page

Virtual Server

Name

Required config name. No spaces.

Status

☐ Disable
 ☒ Enable
 ☐ Maintain

Type

☐ Layer 7
 ☒ Layer 4
 ☐ Layer 2

Address Type

☒ IPv4
 ☐ IPv6

Traffic Group

default ▼

Specifics

Content Routing

☐ Enable

Packet Forwarding Method

DNAT ▼

General

Configuration

Address

0.0.0.0

Example: 192.0.2.1

Port

80

Default: 80 Range: 0 or 1-65535. You can enter up to 8 numbers or number ranges, e.g., 80-90 100.

Connection Limit

0

Default: 10000 Range: 0-100000000 concurrent connections

Connection Rate Limit

0

Default: 0 (disabled) Range: 0-86400 connections per second

Interface

port1 ▼

Resources

Profile

LB_PROF_TCP ▼

Persistence

Click to select. ▼

Method

LB_METHOD_ROUND_ROBIN ▼

Real Server Pool

test-1 ▼

Traffic Log

Traffic Log


☐ Enable

Comments

Please input comments

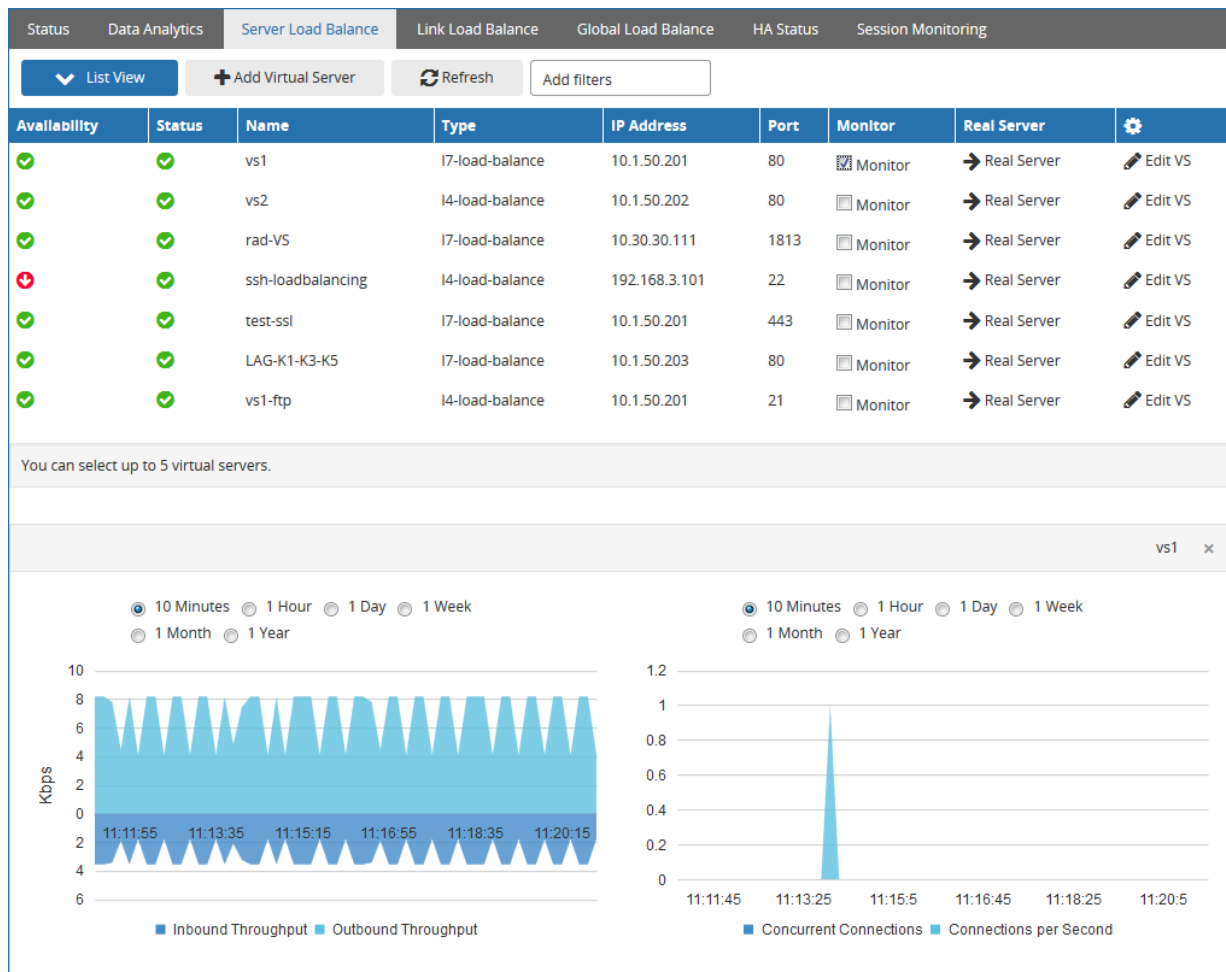
Step 7: Test the deployment

You can test the load balancing deployment by emulating the traffic flow of your planned production deployment. [Figure 12](#) shows a basic network topology.

Figure 12: Basic network topology**To test basic load balancing:**

1. Send multiple client requests to the virtual server IP address.
2. Go to the dashboard to watch the dashboard session and throughput counters increment.
3. Go to Log & Report > Log Browsing > Event Log > Health Check to view health check results.
4. Go to Log & Report > Log Browsing > Traffic Log > SLB HTTP (for example) to view traffic log. It includes throughput per destination IP address.
5. Go to Log & Report > Report to view reports. It has graphs of top N policies and servers.

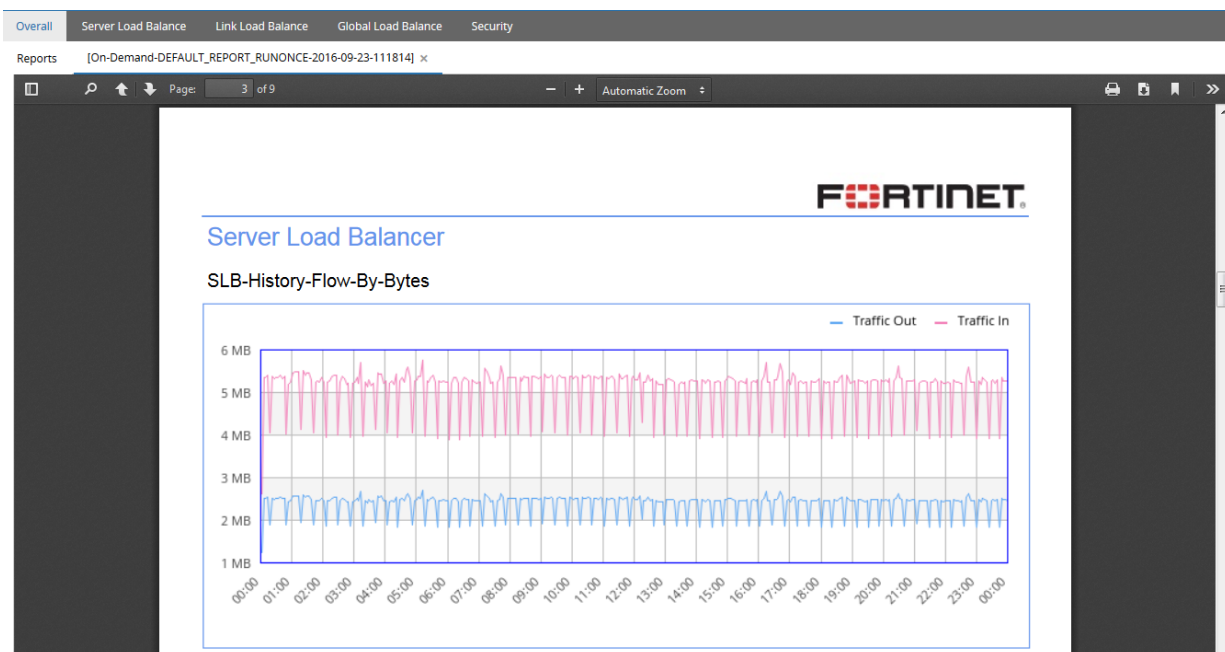
Figure 13 through Figure 16 are examples of the logs and reports you can use to verify your deployment.

Figure 13: Dashboard report**Figure 14: Event log**

Event Log									
<div> Configuration System Admin User Health Check SLB LLB GLB Firewall </div>									
Filter Setting								Download	Refresh
Date	Time	Log Level	Module	Policy	Group	Member	Status	Message	
2016-10-03	11:10:27	alert	slb	vs1-ftp	none	none	success	Virtual server vs1-ftp, status is healthy	📄
2016-10-03	11:10:27	alert	slb	LAG-K1-K3-K5	none	none	success	Virtual server LAG-K1-K3-K5, status is healthy	📄
2016-10-03	11:10:26	alert	slb	test-ssl	none	none	success	Virtual server test-ssl, status is healthy	📄
2016-10-03	11:10:26	alert	slb	ssh-loadbalancing	none	none	failure	Virtual server ssh-loadbalancing, status is down	📄
2016-10-03	11:10:26	alert	slb	rad-VS	none	none	success	Virtual server rad-VS, status is healthy	📄
2016-10-03	11:10:26	alert	slb	vs2	none	none	success	Virtual server vs2, status is healthy	📄
2016-10-03	11:10:26	alert	slb	vs1	none	none	success	Virtual server vs1, status is healthy	📄

Figure 15: Traffic log

Event Log Security Log Traffic Log Script Log												
<input type="radio"/> SLB Layer 4 <input checked="" type="radio"/> SLB HTTP <input type="radio"/> SLB TCPS <input type="radio"/> SLB RADIUS <input type="radio"/> GLB <input type="radio"/> SLB SIP <input type="radio"/> SLB RDP <input type="radio"/> SLB DNS												
Filter Setting Download Refresh												
Date	Time	Source	Received Bytes	Destination	Sent Bytes	Service	Method	URL	Return Code	Virtual Server	Real Server Name	
2016-10-12	11:26:21	10.1.50.101	163	10.1.50.201	461	http	get	/	200	vs1	10_1_51_101	
2016-10-12	11:26:21	10.1.50.101	161	10.1.50.201	461	http	get	/	200	vs1	10_1_51_102	
2016-10-12	11:26:21	10.1.50.101	180	10.1.50.201	461	http	get	/index.html	200	vs1	10_1_51_101	

Figure 16: Overall report

Step 8: Back up the configuration

Once you have tested your basic installation and verified that it functions correctly, create a backup. This “clean” backup is a reference point that has many benefits, including:

- Troubleshooting—You can use a diff tool to compare a problematic configuration with this baseline configuration.
- Restarting—You can rapidly restore your system to a simple yet working point.
- Rapid deployment—You can use the configuration file as a template for other FortiADC systems. You can edit use any text editor to edit the plain text configuration file and import it into another FortiADC system. You should change unique identifiers, such as IP address and sometimes other local network settings that differ from one deployment to another.

To backup the system configuration:

1. Go to System > Settings.
2. Click the **Backup & Restore** tab to display the backup and restore page.
3. Click **Back Up**.

For detailed information, refer to the online help or see [Backing up and restoring the configuration](#).

Chapter 4: Server Load Balancing

This chapter includes the following topics:

- [Server load balancing basics](#)
- [Server load balancing configuration overview](#)
- [Configuring real server SSL profiles](#)
- ["Configuring client SSL profiles" on page 69](#)
- [Configuring real server pools](#)
- [Configuring persistence rules](#)
- [Configuring content routes](#)
- [Using content rewriting rules](#)
- [Configuring compression rules](#)
- [Using caching features](#)
- [Configuring Application profiles](#)
- ["Configuring MySQL profiles" on page 61](#)
- [Configuring error pages](#)
- [Using source pools](#)
- [Configuring authentication policies](#)
- ["Configuring load-balancing \(LB\) methods" on page 152](#)
- [Configuring an L2 exception list](#)
- [Using the Web Category tab](#)
- [Creating a Web Filter Profile configuration](#)
- [Configuring virtual servers](#)

Server load balancing basics

An application delivery controller (ADC) is like an advanced server load balancer. An ADC routes traffic to available destination servers based on health checks and load-balancing algorithms. ADCs improve application availability and performance, which directly improves user experience.

The physical distance between clients and the servers in your backend server farm has a significant impact on server response times. Besides physical distance, the most important factors contributing to server performance are:

- Number of simultaneous connections and requests that the servers can handle
- Load distribution among the servers


The purpose of an ADC is to give you multiple methods for optimizing server response times and server capacity.

After you have deployed an ADC, traffic is routed to the ADC *virtual server* instead of the destination *real servers*.

Figure 17 shows an example of a basic load balancing deployment. The FortiADC appliance is deployed in front of a server farm, and the network interfaces are connected to three subnets: a subnet for management traffic; a

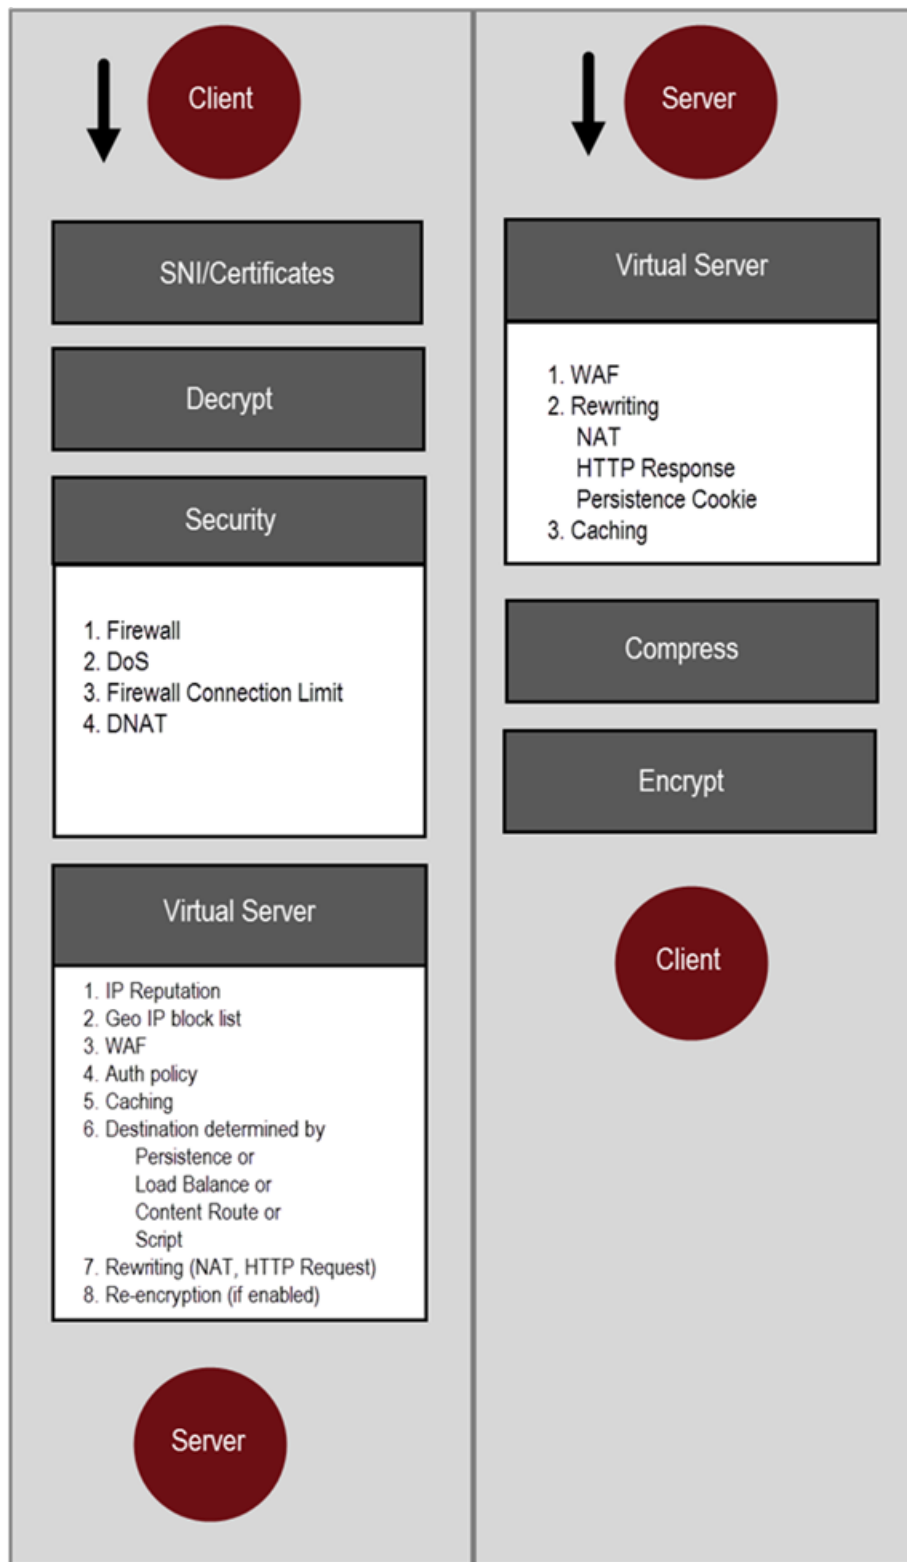
subnet that hosts real servers A, B, and C; and a different subnet that hosts real servers D, E, and F. The FortiADC system performs health checks on the real servers and distributes traffic to them based on system logic and user-defined settings.

Figure 17: Basic network topology



Optionally, you can further improve application security and performance by offloading system processes from the server and having them handled transparently by the ADC. Server tasks that can be handled by the FortiADC appliance include SSL encryption/decryption, WAF protection, Gzip compression, and routing processes, such as NAT.

Figure 18 shows the order in which the FortiADC features process client-to-server and server-to-client traffic.

Figure 18: FortiADC processing

In the client-to-server direction:

- If SNI or SSL decryption is applicable, the system acts on those exchanges.
- Then, security module rules filter traffic, and traffic not dropped continues to the virtual server module.
- Virtual server security features are applied. Traffic not dropped continues for further processing.
- If a caching rule applies, the FortiADC cache serves the content and the request is not forwarded to a backend server.
- If the system selects a destination server based on a persistence rule, content route, or script, the load balancing rules are not applied.
- After selecting a server, the system performs any rewriting and re-encryption actions that are applicable, and then forwards the packets to the server.

In the server-to-client direction:

- WAF HTTP response, NAT, rewriting, persistence, and caching rules are applied.
- If applicable, the FortiADC compresses and encrypts the server response traffic.

Server load balancing configuration overview

The configuration object framework supports the granularity of FortiADC application delivery control rules. You can configure specific options and rules for one particular type of traffic, and different options and rules for another type.

Figure 19 shows the configuration objects used in the server load balancing configuration and the order in which you create them.

Basic steps

1. Configure health check rules and real server SSL profiles.
This step is optional. In many cases, you can use predefined health check rules and predefined real server SSL profiles. If you want to use custom rules, configure them before you configure the pools of real servers.
2. Configure server pools.
This step is required. Server pools are the backend servers you want to load balance and specify the health checks used to determine server availability.
3. Configure persistence rules, optional features and policies, profile components, and load balancing methods.
You can skip this step if you want to select from predefined persistence rules, profiles, and methods.
4. Configure the virtual server.
When you configure a virtual server, you select from predefined and custom configuration objects.

Example workflow

For a members-only HTTPS web server farm, you might have a workflow similar to the following:

1. Configure security module firewall rules that allow only HTTPS traffic from untrusted subnets to the virtual server.
2. Import server SSL certificates, configure a local certificate group, and a certificate verification policy.
3. Configure HTTPS health checks to test the availability of the web servers.
4. Configure the server pools, referencing the health check configuration object.
5. Configure authentication:

- Create a RADIUS or LDAP server configuration.
 - Create user groups.
 - Create an authentication policy.
6. Configure an HTTPS profile, referencing the certificate group and certificate verification policy and setting SSL version and cipher requirements.
 7. Configure an application profile and client SSL profile if needed.
 8. Configure the virtual server, using a combination of predefined and user-defined configuration objects:
 - Predefined: WAF policy, Persistence, Method
 - User-defined: Authentication Policy, Profile

Figure 19: Server load balancing configuration steps

Configuring real server SSL profiles

A real server SSL profile determines settings used in network communication on the FortiADC-server segment, in contrast to a virtual server profile, which determines the settings used in network communication on the client-FortiADC segment.

Figure 20 illustrates the basic idea of client-side and server-side profiles.

Figure 20: SSL profiles

Table 2 provides a summary of the predefined profiles. You can select predefined profiles in the real server pool configuration, or you can create user-defined profiles.

Table 2: Predefined real server profiles

Profile	Defaults
LB_RS_SSL_PROF_DEFAULT	<ul style="list-style-type: none"> • Allow version: SSLv3, TLSv1.0, TLSv1.1, and TLSv1.2 • Cipher suite list.
LB_RS_SSL_PROF_ECDSA	<ul style="list-style-type: none"> • Allow version: SSLv3, TLSv1.0, TLSv1.1, and TLSv1.2 • Cipher suite list: ECDHE-ECDSA-AES256-GCM-SHA384, ECDHE-ECDSA-AES256-SHA384, ECDHE-ECDSA-AES256-SHA, ECDHE-ECDSA-AES128-GCM-SHA256, ECDHE-ECDSA-AES128-SHA256, ECDHE-ECDSA-AES128-SHA, ECDHE-ECDSA-RC4-SHA, ECDHE-ECDSA-DES-CBC3-SHA
LB_RS_SSL_PROF_ECDSA_SSLV3	<ul style="list-style-type: none"> • Allow version: SSLv3 • Cipher suite list: ECDHE-ECDSA-AES256-SHA, ECDHE-ECDSA-AES128-SHA, ECDHE-ECDSA-RC4-SHA, ECDHE-ECDSA-DES-CBC3-SHA

Profile	Defaults
LB_RS_SSL_PROF_ECDSA_TLS12	<ul style="list-style-type: none"> • Allow version: TLSv1.2 • Cipher suite list: ECDHE-ECDSA-AES256-GCM-SHA384, ECDHE-ECDSA-AES256-SHA384, ECDHE-ECDSA-AES128-GCM-SHA256, ECDHE-ECDSA-AES128-SHA256
LB_RS_SSL_PROF_ENULL	<ul style="list-style-type: none"> • Allow version: SSLv3, TLSv1.0, TLSv1.1, and TLSv1.2 • Cipher suite list: eNull <p>Recommended for Microsoft Direct Access servers where the application data is already encrypted and no more encryption is needed.</p>
LB_RS_SSL_PROF_HIGH	<ul style="list-style-type: none"> • Allow version TLSv1.2 • Cipher suite list: ECDHE-RSA-AES256-GCM-SHA384 ECDHE-RSA-AES256-SHA384 ECDHE-RSA-AES256-SHA DHE-RSA-AES256-GCM-SHA384 DHE-RSA-AES256-SHA256 AES256-GCM-SHA384 AES256-SHA256
LB_RS_SSL_PROF_LOW_SSLV2	<ul style="list-style-type: none"> • Allow version: SSLv2 • Cipher suite list: RC4-MD5
LB_RS_SSL_PROF_LOW_SSLV3	<ul style="list-style-type: none"> • Allow version SSLv3 • Cipher suite list.
LB_RS_SSL_PROF_MEDIUM	<ul style="list-style-type: none"> • Allow version: TLSv1.0, TLSv1.1, and TLSv1.2 • Cipher suite list.
NONE	<ul style="list-style-type: none"> • SSL is disabled.

Before you begin:

- You must have Read-Write permission for Load Balance settings.

To configure custom real server profiles:

1. Go to Server Load Balance > Real Server Pool.
2. Click the **Real Server SSL Profile** tab.
3. Click **Add** to display the configuration editor.
4. Complete the configuration as described in [Table 3](#).
5. Save the configuration.



You can clone a predefined configuration object to help you get started with a user-defined configuration.


To clone a configuration object, click the clone icon  that appears in the tools column on the configuration summary page.

Table 3: Real Server SSL Profile configuration guidelines

Settings	Guidelines
Name	<p>Configuration name. Valid characters are A-Z, a-z, 0-9, _, and -. No spaces. You reference this name in the real server pool configuration.</p> <p>Note: After you initially save the configuration, you cannot edit the name.</p>
SSL	<p>Enable/disable SSL for the connection between FortiADC and the real server.</p> <p>Note: The following fields become available only when SSL is enabled. See above.</p>
Customized SSL Ciphers Flag	<p>Enable/disable use of user-specified cipher suites. When enabled, you must select a Customized SSL Cipher. See below.</p>
Customized SSL Ciphers	<p>If the customize cipher flag is enabled, specify a colon-separated, ordered list of cipher suites.</p> <p>An empty string is allowed. If empty, the default cipher suite list is used.</p> <p>The names you enter are validated against the form of the cipher suite short names published on the OpenSSL website:</p> <p>https://www.openssl.org/docs/manmaster/apps/ciphers.html</p>

Settings	Guidelines
New SSL Ciphers	<p>Ciphers are listed from strongest to weakest:</p> <ul style="list-style-type: none"> • ECDHE-ECDSA-AES256-GCM-SHA384 • ECDHE-ECDSA-AES256-SHA384 • ECDHE-ECDSA-AES256-SHA • ECDHE-ECDSA-AES128-GCM-SHA256 • ECDHE-ECDSA-AES128-SHA256 • ECDHE-ECDSA-AES128-SHA • ECDHE-ECDSA-DES-CBC3-SHA • ECDHE-ECDSA-RC4-SHA • ECDHE-RSA-AES256-GCM-SHA384 • ECDHE-RSA-AES256-SHA384 • ECDHE-RSA-AES256-SHA • DHE-RSA-AES256-GCM-SHA384 • DHE-RSA-AES256-SHA256 • DHE-RSA-AES256-SHA • AES256-GCM-SHA384 • AES256-SHA256 • AES256-SHA • ECDHE-RSA-AES128-GCM-SHA256 • ECDHE-RSA-AES128-SHA256 • ECDHE-RSA-AES128-SHA • DHE-RSA-AES128-GCM-SHA256 • DHE-RSA-AES128-SHA256 • DHE-RSA-AES128-SHA • AES128-GCM-SHA256 • AES128-SHA256 • AES128-SHA • ECDHE-RSA-RC4-SHA • RC4-SHA • RC4-MD5 • ECDHE-RSA-DES-CBC3-SHA • EDH-RSA-DES-CBC3-SHA • DES-CBC3-SHA • EDH-RSA-DES-CBC-SHA • DES-CBC-SHA • eNULL <p>We recommend retaining the default list. If necessary, you can deselect ciphers you do not want to support.</p>

Settings	Guidelines
Allow SSL Versions	Select SSL versions that are allowed for the connection.
Certificate Verify	Specify a Certificate Verify configuration object to validate server certificates. This Certificate Verify object must include a CA group and may include OCSP and CRL checks.
SNI Forward Flag	Enable/disable forwarding the client SNI value to the server. The SNI value will be forwarded to the real server only when the client-side ClientHello message contains a valid SNI value; otherwise, nothing is forwarded.
Session Reuse Flag	Enable/disable SSL session reuse.
Session Reuse Limit	The default is 0 (disabled). The valid range is 0-1048576.
TLS Ticket Flag	Enable/disable TLS ticket-based session reuse .

Configuring MySQL profiles

FortiADC (Version 4.7.0 and later) supports MySQL server load-balancing .

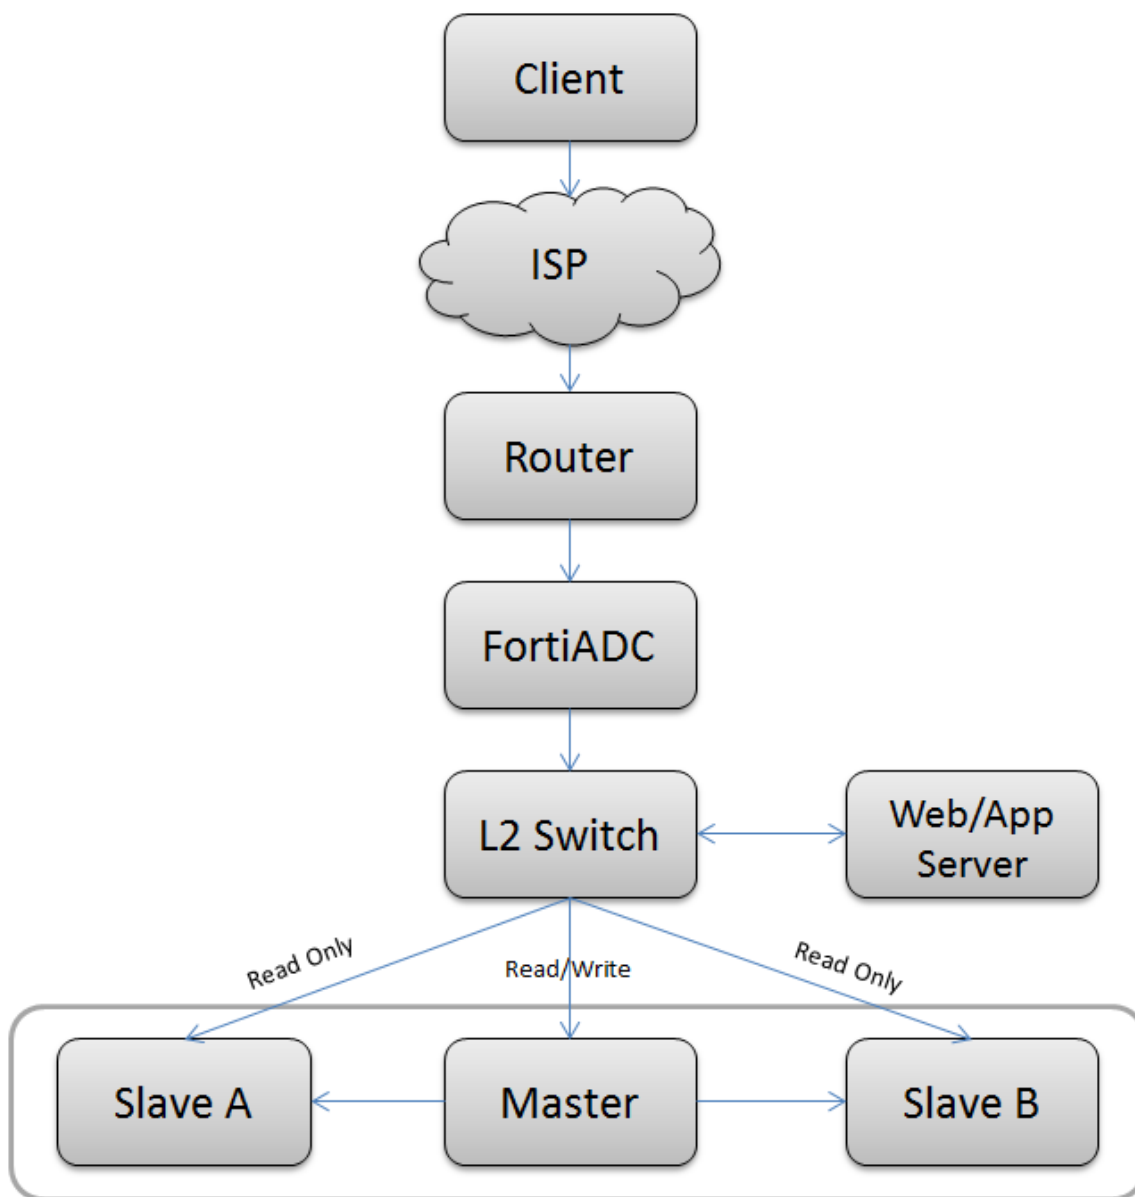
MySQL application profiles are user-specific and must be configured only by the user on a case by case basis. For this reason, FortiADC does not provide any default predefined MySQL application profiles that you can use out of the box. So you must configure your own MySQL load-balancing application profiles to take advantage of this feature.

FortiADC supports two MySQL database load-balancing modes: single master and data sharding.

Single-master mode

The single-master mode is a database server configuration in which a single master MySQL server is responsible for all write operations (i.e., create, update, or delete requests), and one or more slave servers handle all read-only operations. The master server replicates data to the slave servers in a close to real-time fashion. This mode can improve database performance to a certain extent by offloading read-intensive operations to slave servers. It is ideal for load-balancing database traffic that involves more read operations.

The diagram below illustrates the network topology of database server load-balancing in single-master mode.



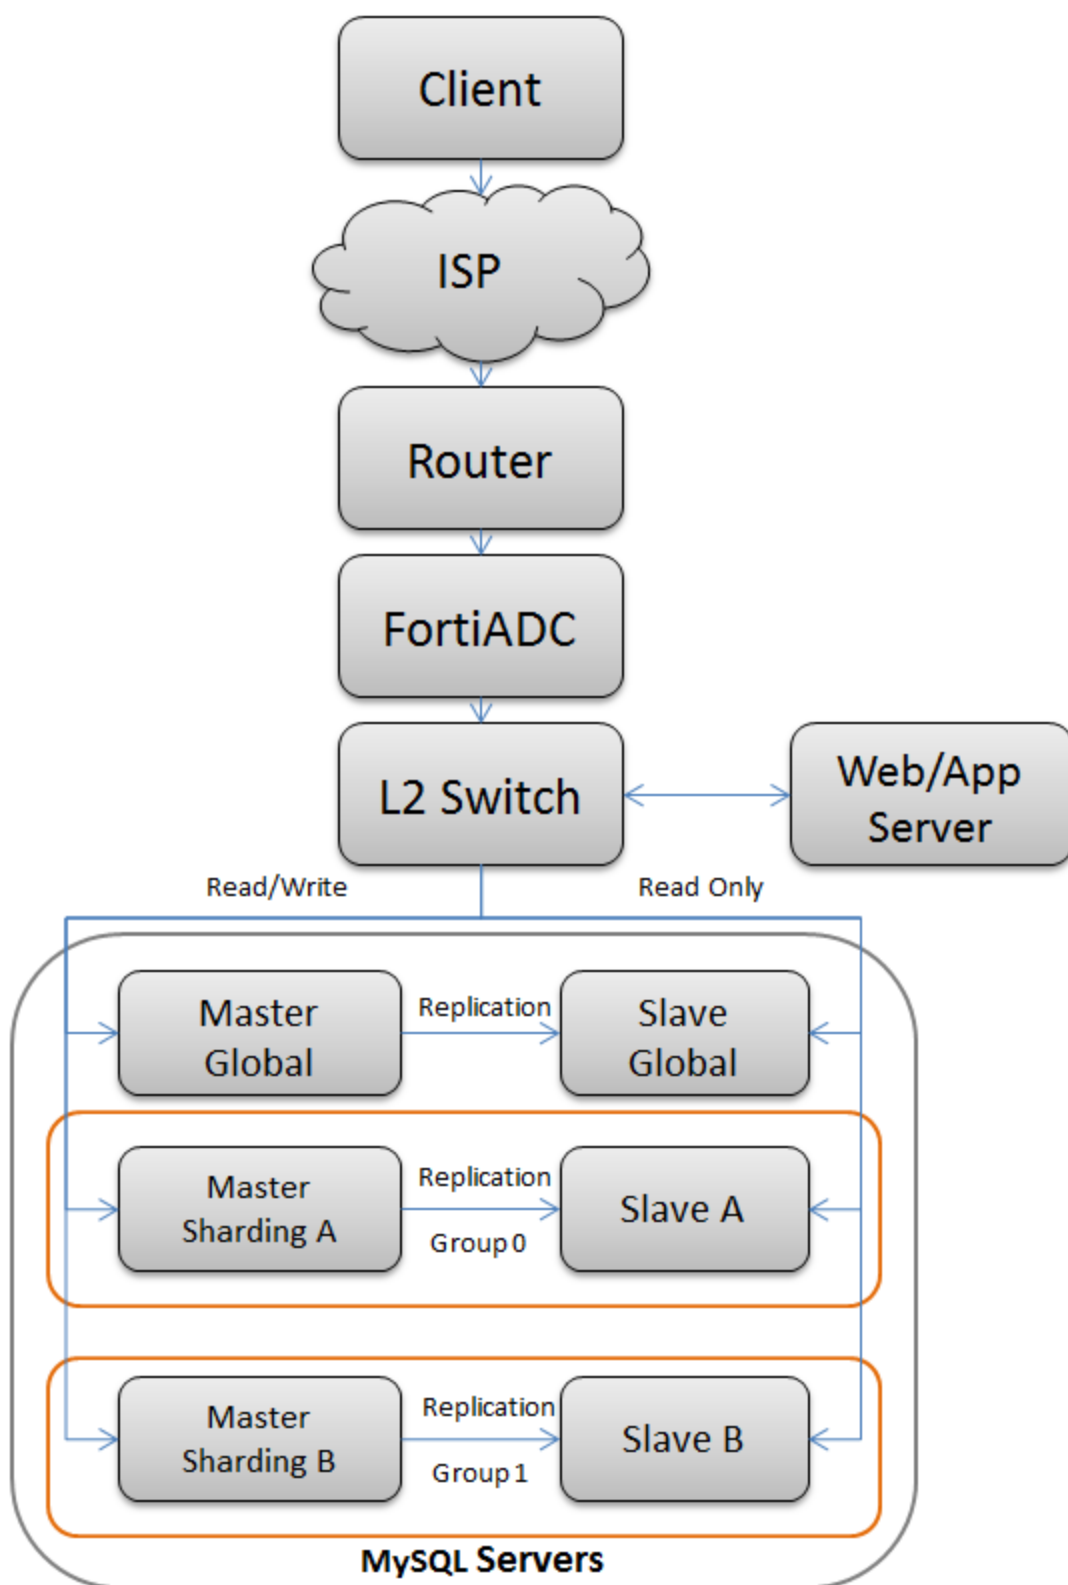
By default, FortiADC passes all write requests to the master server and all read requests (such as select) to the slave servers. So once you have created a MySQL server load-balancing profile, FortiADC will automatically apply this default mode when load-balancing MySQL traffic on the network. However, if you do not like the default behavior, you can change it by setting up your own MySQL server load-balancing rules when configuring your MySQL application profile. For more information, see [Configuring MySQL rules on page 67](#).

Sharding mode

Database sharding is a "shared-nothing" database partitioning technique that breaks down a large database involving a number of database servers into small database chunks and spread them across a number of

distributed servers. It's a highly scalable approach to improving the throughput and performance of large enterprise business applications that are transaction-extensive and database-centric because it provides scalability across independent servers, each having its own CPU, memory, and disks.

The diagram below illustrates MySQL server load-balancing in data-sharding mode.



In sharding mode, FortiADC stores global data on the Master Global—it sends all requests that do not belong to any group to global servers. Using the keys that you have specified, it sends part of the requests to Group 2 and some to Group 1. It supports split read/write in every group.

It must be noted that Data Manipulation Language (DDL) is not supported in sharding mode.

Creating a MySQL profile

Creating a MySQL profile involves the following steps:

1. Create a MySQL configuration object.
2. Specify the existing user name and password of the MySQL database to be used by the MySQL profile configuration object.
3. Configure MySQL Rule (for single-master mode, optional) or MySQL Sharding (for database sharding mode).

Note: You can create MySQL profiles from either the GUI or the CLI. The following paragraphs discuss how to configure a MySQL profile using the GUI. For instructions on how to create MySQL profiles from the CLI, refer to the [FortiADC 4.7.0 CLI Reference](#).

Before you begin:

- You must have already created MySQL database objects to be used by the MySQL profile.
- You must have read-write permission for load-balance settings.

Creating a MySQL configuration object

1. Go to Server Load Balance > Application Resources.
2. Select the **Application Profile** tab if it is not already selected.
3. Click **Add** to open the Application Profile configuration editor.
4. In the **Name** field, enter a unique profile name.
5. In the **Type** field, click the down arrow and select **MySQL** from the drop-down menu.
6. For **MySQL Mode**, select **Single Master** or **Sharding**. Refer to
7. Click **Save**. Your newly created MySQL profile configuration object is automatically appended to the bottom of the Server Load Balancing > Application Resources > Application Profile page.
8. Click the newly created MySQL profile to open it. See the illustration below.

The screenshot shows the 'Application Profile' configuration window. The 'Name' field is 'jack-test-1' and the 'Type' is 'MySQL'. Under 'Specifics', 'MySQL Mode' is set to 'Sharding'. The 'MySQL User Password' section contains a table with two users: 'jack_1' and 'root'. The 'MySQL Sharding' section is empty. The 'MySQL Rule' section is also empty. At the bottom are 'Save' and 'Cancel' buttons.

Application Profile

Application Profile

Name: jack-test-1

Type: MySQL

Specifics

MySQL Mode: ☐ Single Master ☒ Sharding

MySQL User Password

+ Add - Delete Total: 2 Refresh

ID	Username	
1	jack_1	[Edit] [Delete] [Copy]
2	root	[Edit] [Delete] [Copy]

MySQL Sharding

+ Add - Delete Total: 0 Refresh

ID	Table	Type	Key	Group List	
----	-------	------	-----	------------	--

MySQL Rule

+ Add - Delete Total: 0 Refresh

ID	Type	Database List	User List	Table List	Client IP List	SQL List	
----	------	---------------	-----------	------------	----------------	----------	--

Save Cancel

Note: The image above shows a sample MySQL profile configuration object named "jack-test-1". Once a MySQL profile is created, you need to specify the MySQL database user account, and create MySQL Rule or Sharding depending on which MySQL mode you choose to use. The following paragraphs discuss the procedures for each of those tasks.

Specifying the MySQL user account

Once a MySQL profile is created, you must specify a MySQL user account to be used with the profile by entering the user name and password of that account.

It's important to note that you are asked to provide the user name and password of an existing MySQL account. *So do not try to create a new user account here.*

To specify a MySQL user account:

1. In the MySQL User Password pane (see the illustration above), click **Add**. The Edit MySQL User Password dialog opens.
2. Enter the user name and password of the MySQL database account,
3. Click **Save**.

Configuring MySQL rules

When configuring a MySQL rule, you first need to decide whether you want FortiADC to send requests to the Master database server or the Slave database server(s). Then you can set a few conditions (rules) to tell FortiADC how to send the requests. It must be noted that all the conditions are of an "OR" relationship.

To configure a MySQL rule:

1. In the MySQL Rule pane, click **Add**. The Application Profile > Edit MySQL Rule dialog opens.
2. Make the desired entries or selections as described in Table 17.
3. Click **Save**.

Configuring sharding

FortiADC supports two types of database-sharding: by range or by hash. In the former case, FortiADC distributes the data to different groups according to the key range. In the latter case, it first hashes the keys and then automatically distributes the data to different groups.

To configure MySQL sharding:

1. In the MySQL Sharding pane, click **Add**. The Application Profile > Edit MySQL Sharding dialog opens.
2. Make the desired entries or selections as described in Table 17.
3. Click **Save**.

Note: When configuring pool members in the CLI to match the real server pool members on the GUI, you can use the `set mysql-group-id` command to set the groups that match the pool members:

```
config load-balance pool
  edit "sharding"
    set real-server-ssl-profile NONE
    config pool_member
      edit 1
        set pool_member_service_port 3306
        set pool_member_cookie rs
        set real-server master
      next
      edit 2
        set pool_member_service_port 3306
        set pool_member_cookie rs2
        set real-server master2
        set mysql-group-id 1
      next
      edit 3
        set pool_member_service_port 3306
        set pool_member_cookie rs3
        set real-server slave
        set mysql-read-only enable
      next
    
```

```

edit 4
    set pool_member_service_port 3306
    set pool_member_cookie rs4
    set real-server slave2
    set mysql-read-only enable
    set mysql-group-id 1
next
end
next
end

```



You can clone a predefined configuration object to help you get started with a user-defined configuration.


To clone a configuration object, click the clone icon  that appears in the tools column on the configuration summary page.

Table 4: MySQL profile configuration guidelines

Parameter	Description
Application Profile	
Name	A unique name for the MySQL profile you are creating.
Type	MySQL
MySQL Mode	<p>Select either of the following:</p> <p>Single Master—If selected, FortiADC will configure the MySQL profile in single-master mode. See Single-master mode.</p> <p>Sharding—If selected, FortiADC will configure the MySQL profile in database-sharding mode. See Sharding mode.</p>
MySQL User Password	
User Name	The user name of the MySQL database.
Password	The password for the MySQL user name you've entered above.
MySQL Rule	

Parameter	Description
Type	Select either of the following: <ul style="list-style-type: none"> • Master—If selected, FortiADC will send all data specified in the MySQL rule to the master MySQL database server. • Slave—If selected, FortiADC will send all data specified in the MySQL rule to the slave MySQL database server.
Database List	A list of up to eight MySQL database names separated by space
User List	A list of up to eight user names separated by space
Table List	A list of up to eight MySQL Database tables separated by space
Client IP List	A list of up to eight FortiADC client IP addresses separated by space
SQL List	A list of up to eight MySQL statements separated by space
Sharding	
Type	Select either of the following: <ul style="list-style-type: none"> • Range—If selected, FortiADC will send data in the data tables to different groups based on the specified range of the keys. • Hash—If selected, FortiADC will perform hash calculations and then automatically send data to different groups.
Database	The database name
Table	The table name
Key	The column name
Group List	A list of up to eight group IDs Note: The group IDs must match the real server pool members.

Configuring client SSL profiles

A client SSL profile is used to manage the SSL session between the client and the proxy. It allows FortiADC to accept and terminate client requests sent via the SSL protocol. The Client SSL page provides the settings for configuring client-side SSL connections, and displays all the client SSL profiles that have been configured on the system.

Before you begin creating a client SSL profile:

- You must have already created configuration objects for certificates, certificate caching, and certificate verify if you want to include them in the profile.
- You must have read-write permission for Load Balance settings.

To configure custom profiles:

1. Go to Server Load Balance > Application Resources.
2. Click the **Client SSL** tab.
3. Click **Add** to display the configuration editor.
4. Complete the configuration as described in .
5. Save the configuration.



You can clone a predefined client SSL profile to help you get started with a user-defined configuration.


To clone a configuration object, click the clone icon  that appears in the tools column on the configuration summary page.

Table 5: Client SSL profile configuration guidelines

Type	Profile Configuration Guidelines
Name	Specify a unique name for the client SSL profile.
Customized SSL Ciphers Flag	Enable or disable the use of user-specified cipher suites. If enabled, you must specify a colon-separated, ordered list of a customized SSL cipher suites. See below.
Customized SSL Ciphers	Available only when the Customized SSL Cipher Flag is enabled (see above). Specify a colon-separated, ordered list of a customized SSL cipher suites. Note: FortiADC will use the default SSL cipher suite if the field is left empty.

Type	Profile Configuration Guidelines
SSL Ciphers	<p>Ciphers are listed from the strongest to the weakest:</p> <ul style="list-style-type: none"> • ECDHE-ECDSA-AES256-GCM-SHA384 • ECDHE-ECDSA-AES256-SHA384 • ECDHE-ECDSA-AES256-SHA • ECDHE-ECDSA-AES128-GCM-SHA256 • ECDHE-ECDSA-AES128-SHA256 • ECDHE-ECDSA-AES128-SHA • ECDHE-ECDSA-DES-CBC3-SHA • ECDHE-ECDSA-RC4-SHA • ECDHE-RSA-AES256-GCM-SHA384 • ECDHE-RSA-AES256-SHA384 • ECDHE-RSA-AES256-SHA • DHE-RSA-AES256-GCM-SHA384 • DHE-RSA-AES256-SHA256 • DHE-RSA-AES256-SHA • AES256-GCM-SHA384 • AES256-SHA256 • AES256-SHA • ECDHE-RSA-AES128-GCM-SHA256 • ECDHE-RSA-AES128-SHA256 • ECDHE-RSA-AES128-SHA • DHE-RSA-AES128-GCM-SHA256 • DHE-RSA-AES128-SHA256 • DHE-RSA-AES128-SHA • AES128-GCM-SHA256 • AES128-SHA256 • AES128-SHA • ECDHE-RSA-RC4-SHA • RC4-SHA • RC4-MD5 • ECDHE-RSA-DES-CBC3-SHA • EDH-RSA-DES-CBC3-SHA • DES-CBC3-SHA • EDH-RSA-DES-CBC-SHA • DES-CBC-SHA • eNULL <p>Note: We recommend retaining the default list. If necessary, you can deselect the SSL ciphers that you do not want to support.</p>

Type	Profile Configuration Guidelines
Allowed SSL Versions	<p>You have the following options:</p> <ul style="list-style-type: none"> • SSLv2 • SSLv3 • TLSv1.0 • TLSv1.1 • TLSv1.2 <p>We recommend retaining the default list. If necessary, you can deselect SSL versions you do not want to support.</p> <p>Note: FortiADC does not support session reuse for SSLv2 at the client side. Instead, a new SSL session is started.</p>
Client Certificate Verify	Select the client certificate verify configuration object.
Client Certificate Forward	Disabled by default. When enabled, you must specify the client certificate forward header. See below.
Client Certificate Forward Header	When Client Certificate Forward is enabled (see above), specify the client certificate forward header.
Forward Proxy	By default, (SSL) Forward Proxy is disabled. When enabled, you'll have to configure additional settings noted below.
Client SNI Required	Require clients to use the TLS server name indication (SNI) extension to include the server hostname in the TLS client hello message. Then, the FortiADC system can select the appropriate local server certificate to present to the client.
Local Certificate Group	Select a local certificate group that includes the certificates this virtual server presents to SSL/TLS clients. This should be the backend servers' certificate, NOT the appliance's GUI web server certificate. See Manage certificates .
Note: The following fields become available only when Forward Proxy is enabled.	
Forward Proxy Certificate Caching	Select a Forward Proxy Certificate Caching rule.
Forward Proxy Local Signing CA	Select a Forward Proxy Local Signing CA.
Forward Proxy Intermediate CA Group	Select a Forward Proxy Intermediate CA Group.

Type	Profile Configuration Guidelines
Backend SSL SNI Forward	Disabled by default. Enable it to let FortiADC forward Server Name Indication (SNI) from the client to the back end.
Backend Customized SSL Ciphers Flag	Enabled by default. In this case, you must specify the backend customized SS ciphers. See below.
Backend Customized SSL Ciphers	Specify the customized SSL ciphers to be supported at the back end.
Backend Allowed SSL Versions	<p>We recommend retaining the default list. If necessary, you can deselect SSL versions you do not want to support.</p> <p>Note: FortiADC does not support session reuse for SSLv2 at the client side. Instead, a new SSL session is started.</p>

Using real server pools

This section includes the following topics:

- [Configuring real server pools](#)
- [Example: Using port ranges and the port 0 configuration](#)

Configuring real server pools

Server pools are groups of real servers that host the applications that you load balance.

To configure a server pool:

1. Create a server pool object.
2. Add members.

Before you begin:

- You must have a good understanding and knowledge of the backend server boot behavior, for example, how many seconds it takes to “warm up” after a restart before it can process traffic.
- You must know the IP address and port of the applications.
- If you want to select user-defined health checks, you must create them before creating the pool configuration. See [Configuring health checks](#).
- If you want to select user-defined real server SSL profiles, you must create them before creating the pool configuration. See [Configuring real server SSL profiles](#).
- You must have Read-Write permission for Load Balance settings.

After you have configured a real server pool, you can select it in the virtual server configuration.

To configure a pool:

1. Go to Server Load Balance > Real Server Pool.
The configuration page displays the Real Server tab.

2. Click **Add** to display the configuration editor.
3. Complete the configuration and add members as described in [Table 6](#).
4. Save the configuration.

Table 6: Real Server Pool configuration guidelines

Settings	Guidelines
Real Server Pool	
Name	<p>Configuration name. Valid characters are A-Z, a-z, 0-9, _, and -. No spaces. You reference this name in the virtual server configuration.</p> <p>Note: After you initially save the configuration, you cannot edit the name.</p>
Address Type	<ul style="list-style-type: none"> • IPv4 • IPv6
Health Check	Enable health checking for the pool. You can override this for individual servers in the pool.
Health Check Relationship	<ul style="list-style-type: none"> • AND—All of the selected health checks must pass for the server to be considered available. • OR—One of the selected health checks must pass for the server to be considered available.
Health Check List	Select one or more health check configuration objects.
Real Server SSL Profile	Select a real server SSL profile. Real server SSL profiles determine settings for communication between FortiADC and the backend real servers. The default is NONE, which is applicable for non-SSL traffic.
Member	
Status	<ul style="list-style-type: none"> • Enable—The server can receive new sessions. • Disable—The server does not receive new sessions and closes any current sessions as soon as possible. • Maintain—The server does not receive new sessions but maintains any current connections.
Real Server	<p>Click the down arrow and select a real server configuration object from the list menu.</p> <p>Note: The name of the selected real server pool member will appear in logs and reports.</p>

Settings	Guidelines
Port	<p>Enter the backend server's listening port (number), as described below:</p> <ul style="list-style-type: none"> • HTTP—80, • HTTPS—443 • FTP—21 • SMTP—25 • DNS—53 • POP3—110 • IMAP4—143 • RADIUS—1812 • SNMP—161 <p>Tip: The system uses Port 0 as a “wildcard” port. When configured to use Port 0, the system uses the destination port from the client request. For example, if you specify 0, and the destination port in the client request is 50000, the traffic will be forwarded to Port 50000.</p>
Weight	<p>Assigns relative preference among members—higher values are more preferred and are assigned connections more frequently. The default is 1. The valid range is 1 to 256.</p> <p>All load balancing methods consider weight. Servers are dispatched requests proportional to their weight, relative to the sum of all weights.</p> <p>The following example shows the effect of weight on Round Robin:</p> <ul style="list-style-type: none"> • Server A, Weight 2; Server B, Weight 1: Requests are sent AABAAB. • Server A, Weight 3; Server B, Weight 2: Requests are sent AABAB. <p>For other methods, weight functions as a tie-breaker. For example, with the Least Connection algorithm, requests are sent to the server with the least connections. If the number of connections is equal, the request is sent to the server with the greater weight. For example:</p> <ul style="list-style-type: none"> • Server A, Weight 1, 1 connection • Server B, Weight 2, 1 connection • The next request is sent to Server B.

Settings	Guidelines
Recover	<p>Seconds to postpone forwarding traffic after downtime, when a health check indicates that this server has become available again. The default is 0 (disabled). The valid range is 1 to 86,400 seconds. After the recovery period elapses, the FortiADC assigns connections at the warm rate.</p> <p>Examples of when the server experiences a recovery and warm-up period:</p> <ul style="list-style-type: none"> • A server is coming back online after the health check monitor detected it was down. • A network service is brought up before other daemons have finished initializing and therefore the server is using more CPU and memory resources than when startup is complete. <p>To avoid connection problems, specify the separate warm-up rate, recovery rate, or both.</p> <p>Tip: During scheduled maintenance, you can also manually apply these limits by setting Status to Maintenance instead of Enable.</p> <p>Note: Not applicable for SIP servers.</p>
Warm Up	<p>If the server cannot initially handle full connection load when it begins to respond to health checks (for example, if it begins to respond when startup is not fully complete), indicate how long to forward traffic at a lesser rate. The default is 0 (disabled). The valid range is 1 to 86,400 seconds.</p> <p>Note: Not applicable for SIP servers.</p>
Warm Rate	<p>Maximum connection rate while the server is starting up. The default is 10 connections per second. The valid range is 1 to 86,400 connections per second.</p> <p>The warm up calibration is useful with servers that have the network service brought up before other daemons have finished initializing. As the servers are brought online, CPU and memory are more utilized than they are during normal operation. For these servers, you define separate rates based on warm-up and recovery behavior. For example, if Warm Up is 5 and Warm Rate is 2, the number of allowed new connections increases at the following rate:</p> <ul style="list-style-type: none"> • 1st second—Total of 2 new connections allowed (0+2). • 2nd second—2 new connections added for a total of 4 new connections allowed (2+2). • 3rd second—2 new connections added for a total of 6 new connections allowed (4+2). • 4th second—2 new connections added for a total of 8 new connections allowed (6+2). • 5th second—2 new connections added for a total of 10 new connections allowed (8+2). <p>Note: Not applicable for SIP servers.</p>

Settings	Guidelines
Connection Limit	<p>Maximum number of concurrent connections to the backend server. The default is 0 (disabled). The valid range is 1 to 1,048,576 concurrent connections.</p> <p>Note: Connection Limit is not supported for FTP or SIP servers.</p>
Connection Rate Limit	<p>Limit the number of new connections per second to this server. The default is 0 (disabled). The valid range is 1 to 86,400 connections per second.</p> <p>In Layer 4 deployments, you can apply a connection rate limit per real server and per virtual server. Both limits are enforced.</p> <p>Note: The connection rate limit applies only when the real servers belong to a Layer 4 virtual server. If you add a real server pool with this setting configured to a Layer 7 virtual server, for example, the setting is ignored.</p> <p>Note: Connection Rate Limit is not supported for FTP or SIP servers.</p>
Cookie	<p>Specify the cookie name to be used when cookie-based Layer 7 session persistence is enabled. The cookie is used to create a FortiADC session ID, which enables the system to forward subsequent related requests to the same backend server.</p> <p>If you do not specify a cookie name, it is set to the pool member server name string.</p> <p>Note: This option is NOT applicable for SIP servers.</p>
MySQL Group ID	Specify the MySQL group ID.
MySQL Read Only	Disabled by default. Select the button to enable it.
Backup	<p>Designate this as a backup server to which FortiADC will direct traffic when the other servers in the pool are down. The backup server receives connections when all the other pool members fail the health check or you have manually disabled them.</p> <p>Note: Not applicable for SIP servers.</p>
Health Check Inherit	When enabled, FortiADC will use the pool's health check settings. If disabled, you must select a health check to use with this individual backend server. See below.
Health Check	<p>Select this option to specify a health check configuration object for this server.</p> <p>Note: This option becomes available only when</p>

Settings	Guidelines
Health Check Relationship	<ul style="list-style-type: none"> • AND—All of the selected health checks must pass for the server to be considered available. • OR—One of the selected health checks must pass for the server to be considered available.
Health Check List	Select one or more health check configuration objects. Shift-click to select multiple objects at the same time.
RS Profile Inherit	Enable to inherit the real server SSL profile from the pool configuration. Disable to specify the real server profile in this member configuration. See below.
RS Profile	<p>If RS Profile Inherit (above) is disabled, you must specify a real server SSL profile. A real server SSL profile determines the settings for communication between FortiADC and backend real server.</p> <p>Note: This option becomes available only when RS Profile Inherit is disabled.</p>

Example: Using port ranges and the port 0 configuration

In some deployments, it is advantageous to support listening port ranges for client requests. For example, data centers or web hosting companies sometimes use port numbers to identify their customers. Client A sends requests to port 50000, client B to port 50001, client C to port 50002, and so on.

To support this scenario:

1. On the real servers, configure the listening ports and port ranges according to your requirements.
2. On the FortiADC, when you configure the real server pool member, specify port 0 for the port. The system handles port 0 as a “wildcard” port. When configured to use port 0, the system uses the destination port from the client request. For example, if you specify 0, and the destination port in the client request is 50000, the traffic is forwarded to port 50000.
3. When you configure the virtual server, specify a listening port and port range. The port range is like an offset. If the specified port is 50000 and the port range is 10, the virtual server listens on ports 50000-50009.

Figure 21 and Figure 22 highlight the key FortiADC configuration elements.

Figure 21: Real server port 0 configuration

The screenshot shows the 'Real Server Pool' configuration window with the 'Edit Member' tab selected. The 'Port' field is highlighted in yellow and set to 0. The 'Status' is set to 'Enable'. The 'Server Name' field is empty with a placeholder 'Required. Specify the servi'. The 'Address' is set to '192.0.2.1'. The 'Weight' is set to '1'. The 'Recover' field is set to '0'. The 'Warm Rate' is set to '100'. The 'Cookie' field is empty with a placeholder 'Please input cookie'. The 'Health Check Inherit' checkbox is checked. The 'RS Profile Inherit' checkbox is unchecked. The 'Warm Up' field is set to '0'. The 'Connection Limit' is set to '0'. The 'Connection Rate Limit' is set to '0'. The 'Backup' checkbox is unchecked. The 'RS Profile' dropdown is set to 'NONE'. The 'Save' and 'Cancel' buttons are at the bottom.

Real Server Pool ✕

Real Server Pool Edit Member

Status ☒ Enable ☐ Disable ☐ Maintain

Server Name

Port
Default: 80 Range: 0-65535

Recover
Default: 0 (disabled) Range: 0-86400 seconds

Warm Rate
Default: 10 Range: 1-86400 connections per second

Cookie

Health Check Inherit ☒ Enable

RS Profile Inherit ☐ Enable

Address
Example: 192.0.2.1

Weight
Default: 1 Range: 1-256

Warm Up
Default: 0 (disabled) Range: 0-86400 seconds

Connection Limit
Default: 0 (disabled) Range: 0-1048576 concurrent connections

Connection Rate Limit
Default: 0 (disabled) Range: 0-86400 connections per second

Backup ☐ Enable

RS Profile

Save Cancel

Figure 22: Virtual server port range configuration

The screenshot displays the 'Virtual Server' configuration page in FortiADC. The 'Name' field is set to 'example-port-range'. The 'Status' is 'Enable', 'Type' is 'Layer 7', 'Address Type' is 'IPv4', and 'Traffic Group' is 'default'. Under the 'Specifics' section, 'Content Routing' and 'Content Rewriting' are both disabled, and the 'Transaction Rate Limit' is set to 0. The 'General' section is expanded, showing 'Configuration' and 'Resources' tabs. In the 'Configuration' tab, the 'Address' is '0.0.0.0', the 'Port' is '5000' (highlighted in yellow), and the 'Connection Limit' is '10' (also highlighted in yellow). The 'Interface' is 'port1'. The 'Resources' tab shows 'Profile' as 'LB_PROF_HTTP', 'Persistence' as 'Click to select', 'Method' as 'LB_METHOD_ROUND_ROBIN', 'Real Server Pool' as 'test-1', 'Auth Policy' as 'Click to select', and 'Scripting' as 'Click to select'.

Note: Ports shown on the Dashboard > Virtual Server > Real Server page are for the configured port, so in this case, port 0. The ports shown in traffic logs are the actual destination port, so in this case, port 50000.

Configuring persistence rules

Persistence rules identify traffic that should not be load balanced, but instead forwarded to the same backend server that has seen requests from that source before. Typically, you configure persistence rules to support server

transactions that depend on an established client-server session, like e-commerce transactions or SIP voice calls.

The system maintains persistence session tables to map client traffic to backend servers based on the session attribute specified by the persistence rule.

The persistence table is evaluated before load balancing rules. If the packets received by the ADC match an entry in the persistence session table, the packets are forwarded to the server that established the connection, and load balancing rules are not applicable.

Most persistence rule types have a timeout. When the time that has elapsed since the system last received a request from the client IP address is greater than the timeout, the system does not use the mapping table to forward the request. Instead, it again selects the server using the method specified in the virtual server configuration. Hash-based rule types have a timeout built into the hash algorithm. For other types, you can specify the timeout.

[Table 7](#) describes the predefined persistence rules. You can get started with these commonly used persistence methods or create custom objects.

Table 7: Predefined persistence rules

Predefined	Description
LB_PERSIS_SIP	Persistence based on source IP address or subnet.
LB_PERSIS_CONSISTENT_SIP	Persistence based on a hash of source IP address.
LB_PERSIS_HASH_SRC_ADDR_PORT	Persistence based on a hash that includes source IP address and port.
LB_PERSIS_HASH_COOKIE	Persistence based on a hash of a session cookie provided by the backend server.
LB_PERSIS_RDP_COOKIE	Persistence based on RDP cookie sent by RDP clients in the initial connection request.
LB_PERSIS_SSL_SESS_ID	Persistence based on the SSL session ID.
LB_PERSIS_SIP_CALL_ID	Persistence based on the SIP call ID.
LB_PERSIS_PASSIVE_COOKIE	Persistence based on a passive cookie generated by the server. FortiADC does not generate or manage the cookie, but only observes it in the HTTP stream, thus the name "passive cookie". Also known as "server cookie".

Before you begin:

- You must have a good understanding and knowledge of the applications that require persistent sessions and the methods that can be used to identify application sessions.
- You must have Read-Write permission for Load Balance settings.

After you have configured a persistence rule, you can select it in the virtual server configuration.

To configure a persistence rule:

1. Go to Server Load Balance > Application Resources.
2. Click the **Persistence** tab.
3. Click **Add** to display the configuration editor.
4. Give the rule a name, select the type, and specify rule settings as described in [Table 8](#).
5. Save the configuration.



You can clone a predefined configuration object to help you get started with a user-defined configuration.


To clone a configuration object, click the clone icon  that appears in the tools column on the configuration summary page.

Table 8: Persistence rule guidelines

Settings	Guidelines
Name	Configuration name. Valid characters are A-Z, a-z, 0-9, _, and -. No spaces. You reference this name in the virtual server configuration. Note: After you initially save the configuration, you cannot edit the name.
Type	Select a persistence type.
Source Address	
Source Address	Persistence is based on source IP address.
Timeout	Timeout for an inactive persistence session table entry. The default is 300 seconds. The valid range is 1-86,400.
Subnet Mask Bits (IPv4)	Number of bits in a subnet mask to specify a network segment that should follow the persistence rule. For example, if IPv4 maskbits is set to 24, and the backend server A responds to a client with the source IP 192.168.1.100, server A also responds to all clients from subnet 192.168.1.0/24.
Subnet Mask Bits (IPv6)	Number of bits in a subnet mask to specify a network segment that should follow the persistence rule.

Settings	Guidelines
Match across servers	<p>Enable so clients continue to access the same backend server through different virtual servers for the duration of a session.</p> <p>For example, a client session with a vSphere 6.0 Platform Services Controller (PSC) has connections on the following ports: 443, 389, 636, 2012, 2014, 2020. A FortiADC deployment to load balance a cluster of vSphere PSCs includes Layer 4 virtual server configurations for each of these ports. To ensure a client's connections for a session go to the same backend real server:</p> <ol style="list-style-type: none"> 1. Create a persistence object based on Source Address affinity and select the Match Across Servers option. 2. Select this persistence object in each of the Layer 4 virtual servers configured to load balance the vSphere PSC pool. 3. Select the same real server pool object in each of the Layer 4 virtual servers configured to load balance the vSphere PSC pool. <p>When these options are enabled, FortiADC dispatches the initial connection to a real server destination (for example, RS1) based on the virtual server's load balancing method, and the persistence object is noted in the connection table. Subsequent connection attempts with the same source IP address to any FortiADC virtual server that has this persistence object and real server pool are dispatched to RS1, as long as the session is active.</p> <p>Note: In the Layer 4 virtual server configuration, you specify a packet forwarding method. You can use Source Address persistence with Match Across Servers with any combination of Direct Routing, DNAT, and Full NAT packet forwarding methods. However, with NAT46 and NAT64 packet forwarding methods, the source address type is different from the real server address type. To use Match Across Servers with NAT46 or NAT64, all virtual servers for the application must be configured with the same packet forwarding method: all NAT46 or all NAT64.</p>
Source Address Hash	
Source Address Hash	Persistence is based on a hash of the IP address of the client making an initial request.
Source Address-Port Hash	
Source Address-Port Hash	Persistence is based on a hash of the IP address and port of an initial client request.
HTTP Header Hash	
HTTP Header Hash	Persistence is based on a hash of the specified header value found in an initial client request.
Keyword	A value found in an HTTP header.

Settings	Guidelines
HTTP Request Hash	
HTTP Request Hash	Persistence is based on a hash of the specified URL parameter in an initial client request.
Keyword	A URL parameter.
Cookie Hash	
Cookie Hash	Persistence is based on a hash of the cookie provided by the backend server.
Persistent Cookie	
Persistent Cookie	Persistence is based on the cookie provided in the backend server response. It forwards subsequent requests with this cookie to the original backend server.
Keyword	Backend server cookie name.
Timeout	Timeout for an inactive persistence session table entry. The default is 300 seconds. The valid range is 1-86,400.
Insert Cookie	
Insert Cookie	<p>Persistence is based on a cookie inserted by the FortiADC system.</p> <p>The system inserts a cookie whose name is the value specified by Keyword and whose value is the real server pool member Cookie value and expiration date (if the client does not already have a cookie).</p> <p>For example, if the value of Keyword is <code>sessid</code> and the real server pool member Cookie value is <code>rs1</code>, FortiADC sends the cookie <code>sessid=rs1 U6iFN</code> to the client, where <code>U6iFN</code> is the expiration date as a base64 encoded string.</p>
Keyword	Specifies the cookie name.
Timeout	Timeout for an inactive persistence session table entry. The default is 300 seconds. The valid range is 1-86,400.
Rewrite cookie	

Settings	Guidelines
Rewrite Cookie	<p>Persistence is based on the cookie provided in the backend server response, but the system rewrites the cookie.</p> <p>The system checks the HTTP response for a <code>Set-Cookie:</code> value that matches the value specified by Keyword. It replaces the keyword value with the real server pool member Cookie value.</p> <p>For example, the value of Keyword in the persistence configuration is <code>sessid</code>. The real server pool member Cookie value is <code>rs1</code>. After an initial client request, the response from the server contains <code>Set-Cookie: sessid=666</code>, which FortiADC changes to <code>Set-Cookie: sessid=rs1</code>. FortiADC uses this rewritten value to forward subsequent requests to the same backend server as the original request.</p>
Keyword	Specifies a <code>Set-Cookie:</code> value to match.
Embedded Cookie	
Embedded Cookie	<p>Persistence is based on the cookie provided in the backend server response.</p> <p>Like Rewrite Cookie, the system checks the HTTP response for a <code>Set-Cookie:</code> value that matches the value specified by Keyword in the persistence configuration. However, it preserves the original value and adds the real server pool member Cookie value and a ~ (tilde) as a prefix.</p> <p>For example, the value of Keyword is <code>sessid</code>. The real server pool member Cookie value is <code>rs1</code>. After an initial client request, the response from the server contains <code>Set-Cookie: sessid=666</code>, which the system changes to <code>Set-Cookie: sessid=rs1~666</code>. It uses this rewritten value to forward subsequent requests to the same backend server as the original request.</p>
Keyword	Specifies a <code>Set-Cookie:</code> value to match.
RADIUS Attribute	
RADIUS Attribute	Persistence is based on a specified RADIUS attribute.
Keyword	<p>RADIUS attribute. Specify the RADIUS attribute number. For example, specify 8 for Framed-IP-Address or 31 for Calling-Station-ID. For a list of RADIUS attribute numbers, see https://tools.ietf.org/html/rfc2865#page-22.</p>
Timeout	Timeout for an inactive persistence session table entry. The default is 300 seconds. The valid range is 1-86,400.
Match across servers	Enable so clients continue to access the same backend server through different virtual servers for the duration of a session.

Settings	Guidelines
RDP Cookie	
RDP Cookie	Persistence based on RDP cookie sent by RDP clients in the initial connection request.
SSL Session ID	
SSL Session ID	Persistence is based on SSL session ID.
Timeout	Timeout for an inactive persistence session table entry. The default is 300 seconds. The valid range is 1-86,400.
SIP Call ID	
SIP Call ID	Persistence is based on SIP Call ID. For SIP services, you can establish persistence using Source Address, Source Address Hash, or SIP caller ID.
Timeout	Timeout for an inactive persistence session table entry. The default is 300 seconds. The valid range is 1-86,400.

Configuring content routes

You can use the content routes configuration to select the backend server pool based on matches to TCP/IP or HTTP header values.

Layer 7 content route rules are based on literal or regular expression matches to the following header values:

- [HTTP Host](#)
- [HTTP Referer](#)
- [HTTP Request URL](#)
- [SNI](#)
- Source IP address

You might want to use Layer 7 content routes to simplify front-end coding of your web pages or to obfuscate the precise server names from clients. For example, you can publish links to a simple URL named `example.com` and use content route rules to direct traffic for requests to `example.com` to a server pool that includes `server1.example.com`, `server2.example.com`, and `server3.example.com`.

Layer 4 content route rules are based on literal or regular expression matches to the following header values:

- Source IP address

Before you begin:

- You must have a good understanding of [HTTP header fields](#).
- You must have a good understanding of Perl-compatible regular expressions ([PCRE](#)) if you want to use them in rule matching.
- You must have Read-Write permission for Load Balance settings.

After you have configured a content routing rule, you can select it in the virtual server configuration.

Note: You can select multiple content routing rules in the virtual server configuration. Rules you add to that configuration are consulted from top to bottom. The first rule to match is applied. If the traffic does not match any of the content routing rule conditions specified in the virtual server configuration, the system behaves unexpectedly. Therefore, it is important that you create a “catch all” rule that has no match conditions. In the virtual server configuration, this rule should be ordered last so it can be used to forward traffic to a default pool.

To configure a content route rule:

1. Go to Server Load Balance > Virtual Server.
2. Click the **Content Routing** tab.
3. Click **Add** to display the configuration editor.
4. Complete the configuration as described in [Table 9](#).
5. Save the configuration.

Table 9: Content routes configuration guidelines

Settings	Guidelines
Name	Configuration name. Valid characters are A-Z, a-z, 0-9, _, and -. No spaces. You reference this name in the virtual server configuration. Note: After you initially save the configuration, you cannot edit the name.
Type	<ul style="list-style-type: none"> • Layer 4 • Layer 7
Real Server	Select a real server pool.
Persistence Inherit	Enable to use the persistence object specified in the virtual server configuration.
Persistence	If not using inheritance, select a session persistence type.
Method Inherit	Enable to use the method specified in the virtual server configuration.
Method	If not using inheritance, select a load balancing method type.
Comments	A string to describe the purpose of the configuration, to help you and other administrators more easily identify its use.
Layer 4 Specifics	
IPv4/Mask	Address/mask notation to match the source IP address in the packet header.
IPv6/Mask	Address/mask notation to match the source IP address in the packet header.
Layer 7 Match Condition	

Settings	Guidelines
Object	<p>Select content matching conditions based on the following parameters:</p> <ul style="list-style-type: none">• HTTP Host Header• HTTP Referer Header• HTTP Request URL• SNI• Source IP Address <p>Note: When you add multiple conditions, FortiADC joins them with an AND operator. For example, if you specify both a HTTP Host Header and HTTP Request URL to match, the rule is a match only for traffic that meets both conditions.</p>
Type	<ul style="list-style-type: none">• String• Regular Expression
Content	<p>Specify the string or PCRE syntax to match the header or IP address.</p> <p>Note: An empty match condition matches any HTTP request.</p>
Reverse	Rule matches if traffic does not match the expression.

Using content rewriting rules

This section includes the following topics:

- [Overview](#)
- [Configuring content rewriting rules](#)
- [Example: Redirecting HTTP to HTTPS](#)
- [Example: Rewriting the HTTP response when using content routing](#)
- [Example: Rewriting the HTTP request and response to mask application details](#)
- [Example: Rewriting the HTTP request to harmonize port numbers](#)

Overview

You might rewrite the HTTP request/response and HTTP headers for various reasons, including the following:

- Redirect HTTP to HTTPS
- External-to-internal URL translation
- Other security reasons

[Table 10](#) summarizes the HTTP header fields that can be rewritten.

Table 10: HTTP header rewriting

Direction	HTTP Header
HTTP Request	<ul style="list-style-type: none">• Host• Referer
HTTP Redirect	Location
HTTP Response	Location

The first line of an HTTP request includes the HTTP method, relative URL, and HTTP version. The next lines are headers that communicate additional information. The following example shows the HTTP request for the URL `http://www.example.com/index.html`:

```
GET /index.html HTTP/1.1
Host: www.example.com
Referer: http://www.google.com
```

The following is an example of an HTTP redirect including the HTTP Location header:

```
HTTP/1.1 302 Found
Location: http://www.iana.org/domains/example/
```

You can use literal strings or regular expressions to match traffic to rules. To match a request URL such as `http://www.example.com/index`, you create two match conditions: one for the Host header `www.example.com` and another for the relative URL that is in the GET line: `/index.html`.

For HTTP redirect rules, you can specify the rewritten location as a literal string or as a regular expression. For all other types or rules, you must specify the complete URL as a literal string.

Configuring content rewriting rules

Before you begin:

- You must have a good understanding of [HTTP header fields](#).
- You must have a good understanding of Perl-compatible regular expressions ([PCRE](#)) if you want to use them in rule matching or rewriting.
- You must have Read-Write permission for Load Balance settings.

After you have configured a content rewriting rule, you can select it in the virtual server configuration.

Note: You can select multiple content rewriting rules in the virtual server configuration. Rules you add to that configuration are consulted from top to bottom. The first to match is applied. If the traffic does not match any of the content rewriting rule conditions, the header is not rewritten.

To configure a content rewriting rule:

1. Go to Server Load Balance > Virtual Server.
2. Click the **Content Rewriting** tab.
3. Click **Add** to display the configuration editor.
4. Complete the configuration as described in [Table 11](#).
5. Save the configuration.

Table 11: Content rewriting rule guidelines

Settings	Guidelines
Name	<p>Configuration name. Valid characters are A-Z, a-z, 0-9, _, and -. No spaces. You reference this name in the virtual server configuration.</p> <p>Note: After you initially save the configuration, you cannot edit the name.</p>
Comments	A string to describe the purpose of the configuration, to help you and other administrators more easily identify its use.
Action Type	Select whether to rewrite the HTTP request or HTTP response.
HTTP Request Rewrite Actions	
Rewrite HTTP Header	<p>Host—Rewrites the Host header by replacing the hostname with the string you specify. For Host rules, specify a replacement domain and/or port.</p> <p>URL—Rewrites the request URL and Host header using the string you specify. For URL rules, specify a URL in one of the following formats:</p> <ul style="list-style-type: none"> • Absolute URL — <code>https://example.com/content/index.html</code> • Relative URL — <code>content/index.html</code> <p>If you specify a relative URL, the host header is not rewritten.</p> <p>Referer—Rewrites the Referer header with the URL you specify. For Referer rules, you must specify an absolute URL.</p> <p>Note: The rewrite string is a literal string. Regular expression syntax is not supported.</p>
Redirect	<p>Sends a redirect with the URL you specify in the HTTP Location header field.</p> <p>For Redirect rules, you must specify an absolute URL. For example: <code>https://example.com/content/index.html</code></p> <p>Note: The rewrite string can be a literal string or a regular expression.</p>
Send 403 forbidden	Sends a 403 Forbidden response instead of forwarding the request.
HTTP Response Rewrite Actions	
Rewrite HTTP Location	<p>Rewrites the Location header field in the server response.</p> <p>For Location rules, you must specify an absolute URL. For example: <code>https://example.com/content/index.html</code></p> <p>Note: The rewrite string is a literal string. Regular expression syntax is not supported.</p>
Match Condition	

Settings	Guidelines
Object	<p>Select content matching conditions based on the following parameters:</p> <ul style="list-style-type: none"> • HTTP Host Header • HTTP Location Header • HTTP Referer Header • HTTP Request URL • Source IP Address <p>Note: When you add multiple conditions, FortiADC joins them with an AND operator. For example, if you specify both a HTTP Host Header and HTTP Request URL to match, the rule is a match only for traffic that meets both conditions.</p>
Type	<ul style="list-style-type: none"> • String • Regular Expression
Content	Specify the string or PCRE syntax to match the header or IP address.
Reverse	Rule matches if traffic does not match the expression.

Example: Redirecting HTTP to HTTPS

You can use the content rewriting feature to send redirects. One common case to use redirects is when the requested resource requires a secure connection, but you accidentally type an HTTP URL instead of an HTTPS URL in the web browser.

For HTTP redirect rules, you can specify the rewritten location as a literal string or regular expression.

[Figure 23](#) shows a redirect rule that matches a literal string and rewrites a literal string. In the match condition table, the rule is set to match traffic that has the Host header domain `example.com` and the relative URL `/resource/index.html` in the HTTP request URL. The redirect action sends a secure URL in the Location header: `https://example.com/resource/index.html`.

Figure 23: Redirecting HTTP to HTTPS (literal string)

The screenshot shows the 'Content Rewriting' configuration window. The 'Name' field is 'test-1'. The 'Action Type' is 'Request'. The 'Action' is 'Redirect'. The 'Redirect' field contains 'https://example.com/resource/index.htm'. The 'Comments' field contains 'comments'. The 'Match Condition' section shows a table with 2 conditions.

+ Add		X Delete		Total: 2		Refresh	
ID	Object	Type	Content	Reverse			
1	HTTP Host Header	String	match	✓			
2	HTTP Host Header	String	match	✗			

At the bottom are 'Save' and 'Cancel' buttons.

Regular expressions are a powerful way of denoting all possible forms of a string. They are very useful when trying to match text that comes in many variations but follows a definite pattern, such as dynamic URLs or web page content.

Figure 24 shows a redirect rule that uses PCRE capture and back reference syntax to create a more general rule than the previous example. This rule sends a redirect for all connections to the same URL but over HTTP. In the match condition table, the first regular expression is `(.*)`. This expression matches any HTTP Host header and stores it as capture 0. The second regular expression is `^/(.*)$`. This expression matches the path in the Request URL (the content after the `/`) and stores it as capture 1. The regular expression for the redirect action uses the back reference syntax `https://$0$1`.

Figure 24: Redirecting HTTP to HTTPS (regular expression)

Content Rewriting

Name test-1

Action Type ☒ Request ☐ Response

Action Redirect

Redirect https://\$0/\$1

Comments http-to-https

Match Condition (Empty Match Condition will match anything)

+ Add		X Delete		Total: 2		Refresh
ID	Object	Type	Content	Reverse		
1	HTTP Host Header	Regular Expression	match	✓		
2	HTTP Request URL	Regular Expression	match	✗		

Save **Cancel**

Table 12 describes commonly used PCRE syntax elements. Table 13 gives examples of useful and relevant expressions that were originally submitted to the FortiGate Cookbook. For a deeper dive, consult a PCRE reference.



Regular expressions can involve very computationally intensive evaluations. For best performance, you should only use regular expressions where necessary, and build them with care.

Table 12: Common PCRE syntax elements

Pattern	Usage	Example
()	Creates a capture group or sub-pattern for back-reference or to denote order of operations.	<p>Text: /url/app/app/mapp</p> <p>Regular expression: (/app)*</p> <p>Matches: /app/app</p> <p>Text: /url?paramA=valueA&paramB=valueB</p> <p>Regular expression: (param)A=(value)A&\0B\1B</p> <p>Matches: paramA=valueA&paramB=valueB</p>
\$0, \$1, \$2, ...	<p>Only \$0, \$1,..., \$9 are supported.</p> <p>A back-reference is a regular expression token such as \$0 or \$1 that refers to whatever part of the text was matched by the capture group in that position within the regular expression.</p> <p>Back-references are used whenever you want the output/interpretation to resemble the original match: they insert a substring of the original matching text. Like other regular expression features, back-references help to ensure that you do not have to maintain a large, cumbersome list of all possible URLs.</p> <p>To invoke a substring, use \$<i>n</i> (0 <= <i>n</i> <= 9), where <i>n</i> is the order of appearance of capture group in the regular expression, from left to right, from outside to inside, then from top to bottom.</p>	<p>Let's say the regular expressions in a condition table have the following capture groups:</p> <p>(a) (b) (c (d)) (e)</p> <p>This syntax results in back-reference variables with the following values:</p> <p>\$0 — a</p> <p>\$1 — b</p> <p>\$2 — cd</p> <p>\$3 — d</p> <p>\$4 — e</p>

Pattern	Usage	Example
\	<p>Escape character.</p> <p>Except, if it is followed by an alphanumeric character, the alphanumeric character is <i>not</i> matched literally as usual. Instead, it is interpreted as a regular expression token. For example, \w matches a word, as defined by the locale.</p> <p>Except, if it is followed by regular expression special character:</p> <pre>*. ^\$?+\(\)\{\}\[\]\</pre> <p>When this is the case, the \ escapes interpretation as a regular expression token, and instead treats the character as a normal letter.</p> <p>For example, \\ matches the \ character.</p>	<p>Text: /url?parameter=value</p> <p>Regular expression: \?param</p> <p>Matches: ?param</p>
.	<p>Matches any single character <i>except</i> \r or \n.</p> <p>Note: If the character is written by combining two Unicode code points, such as à where the core letter is encoded separately from the accent mark, this will <i>not</i> match the entire character: it will only match one of the code points.</p>	<p>Text: My cat catches things.</p> <p>Regular expression: c.t</p> <p>Matches: cat cat</p>
+	<p>Repeatedly matches the previous character or capture group, 1 or more times, as many times as possible (also called “greedy” matching) <i>unless</i> followed by a question mark (?), which makes it optional.</p> <p>Does not match if there is not at least 1 instance.</p>	<p>Text: www.example.com</p> <p>Regular expression: w+</p> <p>Matches: www</p> <p>Would also match “w”, “ww”, “www”, or any number of uninterrupted repetitions of the character “w”.</p>

Pattern	Usage	Example
*	<p>Repeatedly matches the previous character or capture group, 0 or more times. Depending on its combination with other special characters, this token could be either:</p> <p>* — Match as <i>many</i> times as possible (also called “greedy” matching).</p> <p>*? — Match as <i>few</i> times as possible (also called “lazy” matching).</p>	<p>Text: www.example.com</p> <p>Regular expression: .*</p> <p>Matches: www.example.com</p> <p>All of any text, except line endings (\r and \n).</p> <p>Text: www.example.com</p> <p>Regular expression: (w)*?</p> <p>Matches: www</p> <p>Would also match common typos where the “w” was repeated too few or too many times, such as “ww” in w.example.com or “www” in www.example.com. It would still match, however, if no amount of “w” existed.</p>
?	<p>Makes the preceding character or capture group optional (also called “lazy” matching).</p> <p>This character has a different significance when followed by =.</p>	<p>Text: www.example.com</p> <p>Regular expression: (www\.)?example.com</p> <p>Matches: www.example.com</p> <p>Would also match example.com.</p>
?=	<p>Looks ahead to see if the next character or capture group matches and evaluate the match based upon them, but does <i>not</i> include those next characters in the returned match string (if any).</p> <p>This can be useful for back-references where you do not want to include permutations of the final few characters, such as matching “cat” when it is part of “cats” but <i>not</i> when it is part of “catch”.</p>	<p>Text: /url?parameter=value&pack</p> <p>Regular expression: p(?=arameter)</p> <p>Matches: p, but only in “parameter”, <i>not</i> in “pack”, which does not end with “arameter”.</p>

Pattern	Usage	Example
^	<p>Matches either:</p> <p>the <i>position</i> of the beginning of a line (or, in multiline mode, the first line), <i>not</i> the first character itself</p> <p>the inverse of a character, but only if ^ is the first character in a character class, such as [^A]</p> <p>This is useful if you want to match a word, but only when it occurs at the start of the line, <i>or</i> when you want to match anything that is <i>not</i> a specific character.</p>	<p>Text: /url?parameter=value</p> <p>Regular expression: ^/url</p> <p>Matches: /url, but <i>only</i> if it is at the beginning of the path string. It will <i>not</i> match "/url" in subdirectories.</p> <p>Text: /url?parameter=value</p> <p>Regular expression: [^u]</p> <p>Matches: /rl?parameter=value</p>
\$	<p>Matches the <i>position</i> of the end of a line (or, in multiline mode, the entire string), <i>not</i> the last character itself.</p>	
[]	<p>Defines a set of characters or capture groups that are acceptable matches.</p> <p>To define a set via a whole range instead of listing every possible match, separate the first and last character in the range with a hyphen.</p> <p>Note: Character ranges are matched according to their numerical code point in the encoding. For example, [0-2] matches any UTF-8 code points from 40 to 42 inclusive: @AB</p>	<p>Text: /url?parameter=value1</p> <p>Regular expression: [012]</p> <p>Matches: 1</p> <p>Would also match 0 or 2.</p> <p>Text: /url?parameter=valueB</p> <p>Regular expression: [A-C]</p> <p>Matches: B</p> <p>Would also match "A" or "C". It would <i>not</i> match "b".</p>
{}	<p>Quantifies the number of times the previous character or capture group may be repeated continuously.</p> <p>To define a varying number repetitions, delimit it with a comma.</p>	<p>Text: 1234567890</p> <p>Regular expression: \d{3}</p> <p>Matches: 123</p> <p>Text: www.example.com</p> <p>Regular expression: w{1,4}</p> <p>Matches: www</p> <p>If the string were a typo such as "ww" or "www", it would also match that.</p>

Pattern	Usage	Example
(?i)	Turns on case-insensitive matching for subsequent evaluation, until it is turned off or the evaluation completes.	Text: /url?Parameter=value Regular expression: (?i)param Matches: Param Would also match pArAM etc.
	Matches <i>either</i> the character/capture group before <i>or</i> after the pipe ().	Text: Host: www.example.com Regular expression: (\r\n)\n\r Matches: The line ending, regardless of platform.

Table 13: PCRE examples submitted to the FortiGate Cookbook

Regular Expression	Usage
[a-zA-Z0-9]	Any alphanumeric character. ASCII only; e.g. does not match é or É.
[#\?](.*)	All parameters that follow a question mark or hash mark in the URL. e.g. #pageView or ?param1=valueA¶m2=valueB...; In this expression, the capture group does not include the question mark or hash mark itself.
\b10\.\.1\.\.1\b	A specific IPv4 address.
\b(25[0-5] 2[0-4][0-9] [01]?[0-9][0-9]?) \.(25[0-5] 2[0-4][0-9] [01]?[0-9][0-9]?) \.(25[0-5] 2[0-4][0-9] [01]?[0-9][0-9]?) \.(25[0-5] 2[0-4][0-9] [01]?[0-9][0-9]?) \b	Any IPv4 address.

Regular Expression	Usage
(?i)\b.*\.(a(c(d e(ro)? f g i m n o q r s (ia)? t y w x z) b(a b d e f g h i(z)? j m n o r s t v w y z) c(a(t)? c d f g h i k l m n o((m)? (op)?) r s u v x y z) d(e j k m o z) e(c d u e g h r s t u) f(i j k m o r) g(a b d e f g h i l m n ov p q r s t u w y) h(k m n r t u) i(d e l m n(fo)?(t)? o q r s t) j(e m o(bs)? p) k(e g h i m n p r w y z) l(a b c i k r s t u vy) m(a c d e g h i l k l m n o(bi)? p q r s t u (seum)? v w x y z) n(a(me)? c e(t)? f g i l o p r u z) o(m rg) p(a e f g h k l m n r(o)? s t w y) qa r(e o s u w) s(a b c d e g h i j k l m n o r s t u v y z) t(c d e f g h j k l m n o p r(avel)? t v w z) u(a g k s y z) v(a c e g i n u) w(f s) xxx y(e t u) z(a m w))\b	Any domain name.
(?i)\bwww\..example\..com\b	A specific domain name.
(?i)\b(.*)\..example\..com\b	Any sub-domain name of example.com.

Example: Rewriting the HTTP response when using content routing

It is standard for web servers to have external and internal domain names. You can use content-based routing to forward HTTP requests to example.com to a server pool that includes server1.example.com, server2.example.com, and server3.example.com. When you use content routing like this, you should also rewrite the Location header in the HTTP response so that the client receives HTTP with example.com in the header and not the internal domain server1.example.com.

Figure 25 shows a content routing rule that maps requests to example.com to a server pool.

Figure 25: Content routing for the example.com pool

The screenshot shows the 'Content Routing' configuration window for a rule named 'example.com'. The 'Type' is set to 'Layer 7'. Under the 'General' tab, the 'Real Server' is 'test-1', 'Persistence' is 'Click to select', and 'Method' is also 'Click to select'. Both 'Persistence' and 'Method' have 'Inherit' checked. The 'Comments' field contains 'external-to-internal-name-map'. The 'Match Condition' section shows a table with one rule: ID 1, Object 'HTTP Host Header', Type 'String', Content 'example.com', and Reverse checked. At the bottom are 'Save' and 'Cancel' buttons.

Content Routing

Name example.com

Type ☐ Layer 4 ☒ Layer 7

General

Real Server test-1

Persistence Click to select ☒ Inherit

Method ☒ Inherit

Comments

external-to-internal-name-map

Match Condition (Empty Match Condition will match anything)

+ Add		X Delete		Total: 1		Refresh	
ID	Object	Type	Content	Reverse			
1	HTTP Host Header	String	example.com	<input checked="" type="checkbox"/>	✎ ✕ 📄		

Save **Cancel**

Figure 26 shows an HTTP response rule that matches a regular expression and rewrites a literal string. In the match condition table, the rule is set to match the regular expression `server.*\example\.com` in the HTTP Location header in the response. The rewrite action specifies the absolute URL `http://www.example.com`.

Figure 26: Rewriting the HTTP response when masking internal server names

Content Rewriting

Content Rewriting

Name c-rewrite-2

Action Type ☐ Request ☒ Response

Action Rewrite HTTP Location

Specifics

Location http://

Comments

comments

Match Condition (Empty Match Condition will match anything)

+ Add		X Delete		Total: 1		Refresh	
ID	Object	Type	Content	Reverse			
1	HTTP Location Header	Regular Expression	server.*example.com	<input checked="" type="checkbox"/>			

Save **Cancel**

Example: Rewriting the HTTP request and response to mask application details

Another use case for external-to-internal URL translation involves masking pathnames that give attackers information about your web applications. For example, the unmasked URL for a blog might be `http://www.example.com/wordpress/?feed=rss2`, which exposes that the blog is a wordpress application. In this case, you want to publish an external URL that does not have clues of the underlying technology. For example, in your web pages, you create links to `http://www.example.com/blog` instead of the backend URL.

On FortiADC, you create two rules: one to rewrite the HTTP request to the backend server and another to rewrite the HTTP response in the return traffic.

Figure 27 shows an HTTP request rule. In the match condition table, the rule is set to match traffic that has the Host header domain `example.com` and the relative URL `/blog` in the HTTP request URL. The rule action rewrites the request URL to the internal URL `http://www.example.com/wordpress/?feed=rss2`.

Figure 27: Rewriting the HTTP request when you mask backend application details

Content Rewriting

Content Rewriting

Name c-rewrite-3

Action Type ☒ Request ☐ Response

Action Rewrite HTTP Header

Specifics

Rewrite Host ☐ Enable

Rewrite URL ☒ Enable

URL Content http://www.example.com/wordpress/?fe

Rewrite Referer ☐ Enable

Comments

comments

Match Condition (Empty Match Condition will match anything)

+ Add		X Delete		Total: 2		Refresh	
ID	Object	Type	Content	Reverse			
2	HTTP Host Header	String	example.com	<input checked="" type="checkbox"/>			
1	HTTP Request URL	String	/blog	<input checked="" type="checkbox"/>			

Save **Cancel**

Figure 28 shows the rule for the return traffic. In the match condition table, the rule is set to match traffic that has the string `http://www.example.com/wordpress/?feed=rss2` in the Location header of the HTTP response. The action replaces that URL with the public URL `http://www.example.com/blog`.

Figure 28: Rewriting the HTTP response when you mask backend application details

Content Rewriting

Name c-rewrite-4

Action Type ☐ Request ☒ Response

Action Rewrite HTTP Location

Specifics

Location http://www.example.com/blog

Comments

comments

Match Condition (Empty Match Condition will match anything)

ID	Object	Type	Content	Reverse	Settings
1	HTTP Location Header	String	http://www.example.com/wordpress/?feed=rss2	<input checked="" type="checkbox"/>	

Save Cancel

Example: Rewriting the HTTP request to harmonize port numbers

The HTTP `Host` header contains the domain name and port. You might want to create a rule to rewrite the port so you can harmonize port numbers that are correlated with your application service. For example, suppose you want to avoid parsing reports on your backend servers that show requests to many HTTP service ports. When you review your aggregated reports, you have records for port 80, port 8080, and so on. You would rather have all HTTP requests served on port 80 and accounted for on port 80. To support this plan, you can rewrite the HTTP request headers so that all the `Host` header in all HTTP requests shows port 80.

Figure 29 shows an HTTP request rule that uses a regular expression to match HTTP `Host` headers for `www.example.com` with any port number and change it to port 80.

Figure 29: Rewriting the HTTP request port number

Content Rewriting

Content Rewriting

Name

c-rewrite-5

Action Type

☒ Request
 ☐ Response

Action

Rewrite HTTP Header

Specifics

Rewrite Host

☒ Enable

Host Content

www.example.com

Rewrite URL

☐ Enable

Rewrite Referer

☐ Enable

Comments

comments

Match Condition (Empty Match Condition will match anything)

+ Add

✕ Delete

Total: 1

Refresh

	ID	Object	Type	Content	Reverse	
<input type="checkbox"/>	1	HTTP Host Header	Regular Expression	www.example.com:.*	<input checked="" type="checkbox"/>	

Save

Cancel

Compression and decompression

FortiADC supports HTTP/HTTPS response compression and request decompression with either gzip or deflate format.

You can offload HTTP/HTTPS response compression to FortiADC to save resources on your back-end servers, and let FortiADC to decompress compressed HTTP/HTTPS client requests for WAF inspection before passing them to your back-end servers.

Configuring compression rules

To offload compression from your back-end servers, you can configure FortiADC to perform HTTP/HTTPS compression on behalf of the server.

The following content types can be compressed:

- application/javascript
- application/soap+xml
- application/x-javascript
- application/xml
- text/css
- text/html
- text/javascript
- text/plain
- text/xml
- custom

Not all HTTP/HTTPS responses should be compressed. Compression offers the greatest performance improvements when applied to URLs whose media types include repetitive text such as tagged HTML and JavaScript. Files that already contain efficient compression such as GIF images usually should not be compressed, as the CPU usage and time spent compressing them will result in an increased delay rather than network throughput improvement. Plain text files where no words are repeated, such as configurations with unique URLs or IPs, also may not be appropriate for compression.

FortiADC supports HTTP/HTTPS response compression in either gzip or deflate format.

Before you begin:

- You must have a good understanding of HTTP/HTTPS compression and knowledge of the content types served from the back-end real servers.
- You must have Read-Write permission for Load Balance settings.

Compression is not enabled by default. After you have configured a compression inclusion rule, you can select it in the profile configuration. To enable compression, select the profile when you configure the virtual server.

To configure a compression rule:

1. Click Server Load Balance > Application Resources.
2. Click the **Compression** tab.
3. Click **Add** to display the configuration editor.
4. Complete the configuration as described in [Table 14](#).
5. Save the configuration.

Table 14: Compression configuration

Settings	Guidelines
Name	<p>Configuration name. Valid characters are A-Z, a-z, 0-9, _, and -. No spaces. You reference this name in the profile configuration.</p> <p>Note: After you initially save the configuration, you cannot edit the name.</p>
URI List Type	<ul style="list-style-type: none"> • Include— Select this option to create a compression inclusion rule. HTTP/HTTPS responses that match the URIs and content types specified in this rule will be compressed by FortiADC before being passed to the client. • Exclude—Select this option to create a compression exclusion rule. HTTP/HTTPS responses that match the URIs and content types specified in this rule will not be compressed by FortiADC before being passed to the client.
URI List	Click Add and specify URIs to build the list. You must use regular expressions beginning with a
Content Types	<p>Click Add and select from the following content types to build the list:</p> <ul style="list-style-type: none"> • application/javascript • application/soap+xml • application/x-javascript • application/xml • text/css • text/html • text/javascript • text/plain • text/xml • custom <p>Note: The "custom" option allows you to specify almost any content/media type, including image files in .JPG, .PNG, and .BMP formats. The default is */*, which means any content/media type.</p>

You can use the CLI to configure advanced options:



```
config load-balance compression
    edit 1
        set cpu-limit {enable | disable}
        set max-cpu-usage <percent> -- max cpu usage for
            compression
        set min-content-length <bytes> -- min bytes for
            compression
    end
```

Configuring decompression rules

If the HTTP request body is compressed, FortiADC cannot pass it to the Web Application Firewall (WAF) which will scan it for potential problems such as a data leak or virus.

To allow FortiADC to pass compressed HTTP client requests to WAF for inspection before forwarding it to your back-end server, you must configure a FortiADC decompression policy.

You can configure FortiADC to temporarily decompress the body of a request based on its file type, which is specified by the HTTP/HTTPS Content-Type: header. The appliance can then inspect the traffic. If no policy-violating content is discovered, it will allow the compressed version of the request to pass to the back-end server.

FortiADC supports HTTP/HTTPS request decompression with either gzip or deflate format. Upon receiving a compressed HTTP/HTTPS request body, FortiADC first extracts the HTTP/HTTPS request body to a temporary buffer and then sends the buffer to the Web Application Firewall (WAF) engine for scanning.

FortiADC supports decompression of the following content-type files:

- application/javascript
- application/soap+xml
- application/x-javascript
- application/xml
- text/css
- text/html
- text/javascript
- text/plain
- text/xml
- custom

Before you begin:

- You must have a good understanding of HTTP decompression and knowledge of the content types served from the backend real servers.
- You must have Read-Write permission for Load Balance settings.

Decompression is not enabled by default. After you have configured a decompression rule, you can select it in the profile configuration. To enable decompression, select the profile when you configure the virtual server.

To configure a decompression rule:

1. Click **Server Load Balance > Application Resources**.
2. Click the **Decompression** tab.
3. Click **Add** to display the configuration editor.
4. Complete the configuration as described in [Table 15](#).
5. Save the configuration.

Table 15: Compression configuration

Settings	Guidelines
Name	<p>Specify a unique name for the decompression rule. Configuration name. Valid characters are A-Z, a-z, 0-9, _, and -. No spaces. You reference this name in the profile configuration.</p> <p>Note: After you initially save the configuration, you cannot edit the name.</p>
URI List Type	<ul style="list-style-type: none"> • Include— Select this option to create a decompression inclusion rule. HTTP/HTTPS responses that match the URIs and content types specified in this rule will be decompressed by FortiADC before being passed to the client. • Exclude—Select this option to create a decompression exclusion rule. HTTP/HTTPS responses that match the URIs and content types specified in this rule will not be decompressed by FortiADC before being passed to the client.
URI List	Click Add and specify URIs to build the list.
Content Types	<p>Click Add and select from the following content types to build the list:</p> <ul style="list-style-type: none"> • application/javascript • application/soap+xml • application/x-javascript • application/xml • text/css • text/html • text/javascript • text/plain • text/xml • custom <p>Note: The "custom" option allows you to specify almost any content/media type, including image files in .JPG, .PNG, and .BMP formats. The default is */*, which means any content/media type.</p>

You can use the CLI to configure decompression rules:



```
config load-balance decompression
  edit <name>
    set cpu-limit {enable | disable}
    set max-cpu-usage [1-100]
    set uri-list-type {include | exclude}
    config uri_list
      edit <ID>
        set uri <refex_pattern>
      next
    end
  config content-types
    edit <ID>
      set content-type <types>
      {
        application/javascript
        application/soap+xml
        application/x-javascript
        application/xml
        custom <plain-string>
        text/css
        text/html
        text/javascript
        text/plain
        text/xml
      }
    next
  end
```

You can use the CLI to select a decompression rule in a server load balance profile (HTTP):

```
config load-balance profile
  edit <name>
    ...
    set decompression <decompression name>
    ...
  next
end
```


Using caching features

The system RAM cache can store HTTP content and serve subsequent HTTP requests for that content without forwarding the requests to the backend servers, thereby reducing the load on the backend servers.

You can configure basic static caching or dynamic caching rules.

Static caching

Figure 30 illustrates the static caching feature.

Figure 30: Static caching feature

Before content is cached	After content has been cached
<ol style="list-style-type: none"> 1. FortiADC receives the request from Client A and checks to see if it has a cached copy of the content. 2. If it does not, it forwards the request to a backend server. 3. The server sends content in response, and FortiADC caches the content. 4. FortiADC sends it to the client. 	<ol style="list-style-type: none"> 1. FortiADC receives the request from Client B and checks to see if it has a cached copy of the content. 2. It does, so it responds by sending the content to the client. The backend server is not contacted.

In general, the RAM cache conforms with the cache requirements described in sections 13 and 14 in [RFC 2616](#).

If caching is enabled for the profile that is applied to traffic processing, the system evaluates HTTP responses to determine whether or not to cache the content. HTTP responses with status codes 200, 203, 300, 301, 400 can be cached.

The following content is not cached:

- A response for a request that uses any method other than GET.
- A response for a request of which URI is contained in URI Exclude List or Dynamic Request URI Invalid list.
- A response for a request that contains any of the following headers: If-Match, If-Unmodified-Since, Authorization, Proxy-Authorization.
- A response that contains any of the following headers: Pragma, Vary, Set-Cookie, and Set-Cookie2.
- A response that does not include the Content-Length header. The Content-Length header must be 0.
- A response that does not contain the following headers: Cache-Control, Expires.
- A response with a Cache-Control header that does not have any of the following values: public, max-age, s-maxage.
- A response with a Cache-Control header that has one of the following values: no-cache, no-store, private.

In addition, content is not cached if the user-configured RAM cache thresholds described below are exceeded.

Dynamic caching

Dynamic caching is subject to rules that you configure. In the Dynamic Caching Rules List, content that matches "caching invalid" URIs is never cached; otherwise, content that matches the Dynamic Cache Rule List of URIs is cached only for the period you specify.

Dynamic caching is useful for dynamic web app experiences, such as online stores. For example, suppose a site uses a shopping cart. The URL to list items in the shopping cart is as follows:

`http://customshop.com/cart/list`

The URLs to add or delete items in the cart is as follows:

`http://customshop.com/cart/add`

`http://customshop.com/cart/delete`

In this case, you never want to cache the added or deleted pages because the old content will be "invalidated" by the changes you make. You may want, however, to cache the list page, but only for the period of time that you specify. The dynamic "invalid" rules makes it possible for you to never cache added and deleted pages, whereas the Dynamic Cache Rule List allows you to cache the list page for a specified period of time.

Another case where dynamic caching is useful is when content on a page is dynamic. For example, suppose an online ticket vendor has the following URL that shows how many tickets remain available for an event: `http://customshop.com/tickets/get_remains`. The number of tickets available is updated by a backend database. In this case, you might want to invalidate caching the URL or specify a shorter age-out time for it.

Configuring caching rules

Before you begin:

- You must have a good understanding of caching and knowledge about the size of content objects clients access on the backend servers.
- You must have deep and detailed knowledge of your website URIs if you want to create dynamic caching rules.
- You must have Read-Write permission for Load Balance settings.

Caching is not enabled by default. After you have configured caching, you can select it in the profile configuration. To enable caching, select the profile when you configure the virtual server.

To configure caching:

1. Click **Server Load Balance > Application Resources**.
2. Click the **Caching** tab.
3. Click **Add** to display the configuration editor.
4. Complete the configuration as described in [Table 16](#).
5. Save the configuration.

Table 16: Caching configuration

Settings	Guidelines
Name	<p>Configuration name. Valid characters are A-Z, a-z, 0-9, _, and -. No spaces. You reference this name in the profile configuration.</p> <p>Note: After you initially save the configuration, you cannot edit the name.</p>
Maximum Object Size	The default is 1 MB. The valid range is 1 byte to 10 MB.
Maximum Cache Size	The default is 100 MB. The valid range is 1 byte to 500 MB.
Maximum Entries	The default is 10,000. The valid range is 1 to 262,144.
Maximum Age	<p>The default is 43,200 seconds. The valid range is 60 to 86,400.</p> <p>The backend real server response header also includes a maximum age value. The FortiADC system enforces whichever value is smaller.</p>
URI Exclude List	
URI	<p>Specify URIs to build the list. You can use regular expressions.</p> <p>This list has precedence over the Dynamic Cache Rule List. In other words, if a URI matches this list, it is ineligible for caching, even if it also matches the Dynamic Cache Rule list.</p>
Dynamic Cache Rule List	
Age	Timeout for the dynamic cache entry. The default is 60 seconds. The valid range is 1-86,400. This age applies instead of any age value in the backend server response header.
URI	<p>Pattern to match the URIs that have content you want cached and served by FortiADC.</p> <p>Be careful with matching patterns and the order rules in the list. Rules are consulted from lowest rule ID to highest. The first rule that matches is applied.</p>
Invalid URI	<p>Pattern to match URIs that trigger cache invalidation.</p> <p>Be careful with matching patterns and the order rules in the list. Rules are consulted from the lowest rule ID to the highest. The first rule that matches is applied.</p> <p>This list has precedence over the Dynamic Cache URI list. In other words, if a URI matches this list, it will not be cached even if it also matches the Dynamic Cache URI list.</p>

Configuring certificate caching

Certificate caching allows the system to cache the certificates presented to it for later use. Once cached, the certificates can be readily retrievable from the cache so that the system does not have to reload them when clients requesting service. In so doing, system performance can be greatly improved.

Configuring a certificate caching object

1. Click Server Load Balance > Application Resources.
2. Click the Certificate Caching tab.
3. Click Add to open the certificate caching editor.
4. Make the desired entries as described in Table 1.
5. Click Save.

Table 17: Certificate caching configuration guidelines

Settings	Guidelines
Name	Enter a unique name for the certificate caching rule.
Maximum Certificate Cache Size	Specify the maximum size of the certificate caching object. The default is 100 M.
Maximum entries	Specify the maximum number of real servers whose certificates (RSA + ECDSA) are to be cached. The default is 100,000.

Configuring Application profiles

An application profile is a configuration object that defines how you want the FortiADC virtual server to handle traffic for specific protocols.

Table 18 describes usage for by application profile type, including compatible virtual server types, load balancing methods, persistence methods, and content route types.

Table 18: Application profile usage

Profile	Usage	VS Type	LB Methods	Persistence
FTP	Use with FTP servers.	Layer 4	Round Robin, Least Connections, Fastest Response	Source Address, Source Address Hash

Profile	Usage	VS Type	LB Methods	Persistence
HTTP	Use for standard, unsecured web server traffic.	Layer 7, Layer 2	Layer 7: Round Robin, Least Connections, URI Hash, Full URI Hash, Host Hash, Host Domain Hash, Dynamic Load Layer 2: Same as Layer 7, plus Destination IP Hash	Source Address, Source Address Hash, Source Address-Port Hash, HTTP Header Hash, HTTP Request Hash, Cookie Hash, Persistent Cookie, Insert Cookie, Embedded Cookie, Rewrite Cookie, Passive Cookie
HTTPS	Use for secured web server traffic when offloading TLS/SSL from the backend servers. You must import the backend server certificates into FortiADC and select them in the HTTPS profile.	Layer 7, Layer 2	Same as HTTP	Same as HTTP, plus SSL Session ID

Profile	Usage	VS Type	LB Methods	Persistence
TURBO HTTP	<p>Use for unsecured HTTP traffic that does not require advanced features like caching, compression, content rewriting, rate limiting, Geo IP blocking, or source NAT. The profile can be used with content routes and destination NAT, but the HTTP request must be in the first data packet.</p> <p>This profile enables packet-based forwarding that reduces network latency and system CPU usage. However, packet-based forwarding for HTTP is advisable only when you do not anticipate dropped packets or out-of-order packets.</p>	Layer 7	Round Robin, Least Connections, Fastest Response	Source Address
RADIUS	Use with RADIUS servers.	Layer 7	Round Robin	RADIUS attribute
RDP	Use with Windows Terminal Service (remote desktop protocol).	Layer 7	Round Robin, Least Connections	Source Address, Source Address Hash, Source Address-Port Hash, RDP Cookie
SIP	Use with applications that use session initiation protocol (SIP), such as VoIP, instant messaging, and video.	Layer 7	Round Robin, URI Hash, Full URI Hash	Source Address, Source Address Hash, Source Address-Port Hash, SIP Call ID

Profile	Usage	VS Type	LB Methods	Persistence
TCP	Use for other TCP protocols.	Layer 4, Layer 2	Layer 4: Round Robin, Least Connections, Fastest Response Layer 2: Round Robin, Least Connections, Fastest Response, Destination IP Hash	Source Address, Source Address Hash
TCPS	Use for secured TCP when offloading TLS/SSL from the backend servers. Like the HTTPS profile, you must import the backend server certificates into FortiADC and select them in the TCPS profile.	Layer 7, Layer 2	Layer 7: Round Robin, Least Connections Layer 2: Round Robin, Least Connections, Destination IP Hash	Source Address, Source Address Hash, Source Address-Port Hash, SSL Session ID
UDP	Use with UDP servers.	Layer 4, Layer 2	Layer 4: Round Robin, Least Connections, Fastest Response, Dynamic Load Layer 2: Same as Layer 4, plus Destination IP Hash	Source Address, Source Address Hash
IP	Combines with Layer 2 TCP/UDP/HTTP virtual server to balance the rest of the IP packets passed through FortiADC. When running the IP protocol 0 VS, the traffic always tries to match none protocol 0 VS first.	Layer 2	Round Robin only.	Source Address, Source Address Hash
DNS	Use with DNS servers.	Layer 7	Round Robin, Least Connections	Not supported yet.
SMTP	Use with SMTP servers.	Layer 7	Round Robin, Least Connections	Source Address, Source Address Hash

Profile	Usage	VS Type	LB Methods	Persistence
RTMP	A TCP-based protocol used for streaming audio, video, and data over the Internet	Layer 7	Round Robin, Least Connection	Source Address, Source Address Hash
RTSP	A network control protocol used for establishing and controlling media sessions between end points	Layer 7	Round Robin, Least Connection	Source Address, Source Address Hash
MySQL	MySQL network protocol stack (i.e., MySQL-Proxy) which parses and builds MySQL protocol packets	Layer 7	Round Robin, Least Connection	N/A

Table 19 shows the default values of the predefined profiles. All values in the predefined profiles are view-only, and cannot be modified. You can select predefined profiles in the virtual server configuration, or you can create user-defined profiles, especially to include configuration objects like certificates, caching settings, compression options, and IP reputation.

Table 19: Predefined profiles

Profile	Defaults
LB_PROF_TCP	Timeout TCP Session—100 Timeout TCP Session after FIN—100 IP Reputation—Disabled Customized SSL Ciphers Flag—Disabled Geo IP block list—None Geo IP Whitelist—None
LB_PROF_UDP	Timeout UDP Session—100 IP Reputation—Disabled Customized SSL Ciphers Flag—Disabled Geo IP block list—None Geo IP Whitelist—None

Profile	Defaults
LB_PROF_HTTP	Client Timeout—50 Server Timeout—50 Connect Timeout—5 Queue Timeout—5 HTTP Request Timeout—50 HTTP Keepalive Timeout—50 Buffer Pool—Enabled Source Address—Disabled X-Forwarded-For—Disabled X-Forwarded-For Header—Blank IP Reputation—Disabled HTTP Mode—Keep Alive Customized SSL Ciphers Flag—Disabled Compression—None. Decompression—None Caching—None Geo IP Block List—None Geo IP Whitelist—None Geo IP Redirect URL—http://

Profile	Defaults
LB_PROF_HTTP_SERVERCLOSE	Client Timeout—50 Server Timeout—50 Connect Timeout—5 Queue Timeout—5 HTTP Request Timeout—50 HTTP Keepalive Timeout—50 Buffer Pool—Enabled Source Address—Disabled X-Forwarded-For—Disabled X-Forwarded-For Header—None IP Reputation—Disabled HTTP Mode—Server Close Customized SSL Ciphers Flag—Disabled Compression—None Decompression—None Caching—None Geo IP Block List—None Geo IP Whitelist—None Geo IP Redirect URL—http://
LB_PROF_TURBOHTTP	Timeout TCP Session—100 Timeout TCP Session after FIN—100 IP Reputation—Disabled Customized SSL Ciphers Flag—Disabled Geo IP Block List—None Geo IP Whitelist—None

Profile	Defaults
LB_PROF_FTP	Timeout TCP Session—100 Timeout TCP Session after FIN—100 IP Reputation—Disabled Customized SSL Ciphers Flag—Disabled Geo IP Block List—None Geo IP Whitelist—None
LB_PROF_RADIUS	Customized SSL Ciphers Flag—Disabled Session Timeout—300 Geo IP Block List—None Geo IP Whitelist—None
LB_PROF_SIP	SIP Max Size—65535 Server Keepalive Timeout—30 Server Keepalive—Enabled Client Keepalive—Disabled Client Protocol—UDP Server Protocol—None Failed Client Type—Drop Failed Server Type—Drop Insert Client IP—Disabled Customized SSL Ciphers Flag—Disabled Geo IP Block List—None Geo IP Whitelist—None

Profile	Defaults
LB_PROF_RDP	Client Timeout—50 Server Timeout—50 Connect Timeout—5 Queue Timeout—5 Buffer Pool—Enabled Source Address—Disabled IP Reputation—Disabled Customized SSL Ciphers Flag—Disabled Geo IP Block List—None Geo IP Whitelist—None
LB_PROF_IP	IP Reputation—Disabled Customized SSL Ciphers Flag—Disabled Geo IP Block List—None Geo IP Whitelist—None Timeout IP Session—100
LB_PROF_DNS	DNS Cache Flag—Enabled DNS Cache Ageout Time—3600 DNS Cache Size—10 DNS Cache Entry Size—512 DNS Cache Response Type—All Records DNS Malform Query Action—Drop DNA Max Query Length—512 DNS Authentication Flag—Disabled

Profile	Defaults
LB_PROF_TCPS	Client Timeout—50 Server Timeout—50 Connect Timeout—5 Queue Timeout—5 Buffer Pool—Enabled Source Address—Disabled IP Reputation—Disabled Customized SSL Ciphers Flag—Disabled SSL Cipher—Shows all available SSL Ciphers, with the default ones selected. Allow SSL Versions—SSLv3, TLSv1.0, TLSv1.1, TLSv1.2 Client SNI Required—Disabled Geo IP block list—None SSL Ciphers—None Client SNI Required—disabled Certificate Group—LOCAL_CERT_GROUP Certificate Verify—None

Profile	Defaults
LB_PROF_HTTPS	<p>Client Timeout—50</p> <p>Server Timeout—50</p> <p>Connect Timeout—5</p> <p>Queue Timeout—5</p> <p>HTTP Request Timeout—50</p> <p>HTTP Keepalive Timeout—50</p> <p>Buffer Pool—Enabled</p> <p>Source Address—Disabled</p> <p>X-Forwarded-For—Disabled</p> <p>X-Forwarded-For Header—None</p> <p>IP Reputation—Disabled</p> <p>HTTP Mode—Keep Alive</p> <p>SSL Proxy Mode—Disabled</p> <p>Customized SSL Ciphers Flag—Disabled</p> <p>SSL Cipher—Shows all available SSL ciphers, with the default ones selected</p> <p>Allow SSL Versions—SSLv3, TLSv1.0, TLS1.1, TLSv1.2</p> <p>Client SNI Required—Disabled</p> <p>Compression—None</p> <p>Decompression—None</p> <p>Caching—None</p> <p>Geo IP Block List—None</p> <p>Geo IP Whitelist—None</p> <p>Geo IP Redirect URL—http://</p> <p>Certificate Group—LOCAL_CERT_GROUP</p> <p>Certificate Verify—None</p>

Profile	Defaults
LB_PROF_HTTPS_SERVERCLOSE	<p>Client Timeout—50</p> <p>Server Timeout—50</p> <p>Connect Timeout—5</p> <p>Queue Timeout—5</p> <p>HTTP Request Timeout—50</p> <p>HTTP Keepalive Timeout—50</p> <p>Buffer Pool—Enabled</p> <p>Source Address—Disabled</p> <p>X-Forwarded-For—Disabled</p> <p>X-Forwarded-For Header—None</p> <p>IP Reputation—Disabled</p> <p>HTTP Mode—Server Close</p> <p>SSL Proxy Mode—Disabled</p> <p>Customized SSL Ciphers Flag—Disabled</p> <p>SSL Cipher—Shows all available SSL ciphers, with the default ones selected</p> <p>Allow SSL Versions—SSLv3, TLSv1.0, TLS1.1, TLSv1.2</p> <p>Client SNI Required—Disabled</p> <p>Compression—None</p> <p>Decompression—None</p> <p>Caching—None</p> <p>Geo IP Block List—None</p> <p>Geo IP Whitelist—None</p> <p>Geo IP Redirect URL—http://</p> <p>Certificate Group—LOCAL_CERT_GROUP</p> <p>Certificate Verify—None</p>

Profile	Defaults
LB_PROF_SMTP	Starttls Active Mode—require Customized SSL Ciphers Flag—Disabled SSL Ciphers—Shows all available SSL Ciphers, with the defaults ones selected Allow SSL Versions —SSLv3, TLSv1.0, TLSv1.1, TLSv1.2 Forbidden Command—expn, turn, vrfy Local Certificate Group—LOCAL_CERT_GROUP
LB_PROF_RTSP	Max Header Size—Default is 4096. Valid values range from 2048 to 65536. Source Address—Disabled by default. When enabled, FortiADC will use the client address to connect to the server pool.
LB_PROF_RTMP	Source Address—Disabled by default. When enabled, FortiADC will use the client address to connect to the server pool.

Before you begin:

- You must have already created configuration objects for certificates, caching, and compression if you want the profile to use them.
- You must have Read-Write permission for Load Balance settings.

To configure custom profiles:

1. Go to Server Load Balance > Application Resources. Click the **Application Profile** tab.
2. Click **Add** to display the configuration editor.
3. Give the profile a name, select a protocol type; then complete the configuration as described in [Table 20](#).
4. Save the configuration.



You can clone a predefined configuration object to help you get started with a user-defined configuration.


To clone a configuration object, click the clone icon  that appears in the tools column on the configuration summary page.

Table 20: Profile configuration guidelines

Type	Profile Configuration Guidelines
TCP	
Timeout TCP Session	Client-side timeout for connections where the client has not sent a FIN signal, but the connection has been idle. The default is 100 seconds. The valid range is 1 to 86,400.
Timeout TCP Session after FIN	Client-side connection timeout. The default is 100 seconds. The valid range is 1 to 86,400.
IP Reputation	Enable to apply the FortiGuard IP reputation service. See Managing IP Reputation policy settings .
Geo IP Block List	Select a Geo IP block list configuration object. See Using the Geo IP block list .
Geo IP Whitelist	Select a whitelist configuration object. See Using the Geo IP whitelist .
IP	
IP Reputation	Enable to apply FortiGuard IP reputation service. IP reputation. See Managing IP Reputation policy settings .
Geo IP Block List	Select a Geo IP block list configuration object. See Using the Geo IP block list .
Geo IP Whitelist	Select a whitelist configuration object. See Using the Geo IP whitelist .
Timeout IP Session	Client-side session timeout. The default is 100 seconds. The valid range is 1 to 86,400.
DNS	
Customized SSL Ciphers Flag	Enable or disable the Customized SSL Ciphers Flag.
DNS Cache Flag	Enable/Disable DNS cache flag.
Geo IP Block List	Select a Geo IP block list configuration object. See Using the Geo IP block list .
Geo IP Whitelist	Select a whitelist configuration object. See Using the Geo IP whitelist .
DNS Cache Flag	Enable or disable the DNS Cache Flag.
DNS Cache Ageout Time	Enter a value from 0 to 65,535. The default is 3,600.

Type	Profile Configuration Guidelines
DNS Cache Size	Enter a value from 1 to 100. The default is 10.
DNS Cache Entry Size	Enter a value from 256 to 4,096. The default is 512.
DNS Malform Query Action	Choose either of the following: <ul style="list-style-type: none"> • Drop • Forward
DNS Max Query Length	Enter a value from 256 to 4,096. The default is 512.
DNS Authentication Flag	Enable or disable DNS authentication flag.
UDP	
Timeout UDP Session	Client-side session timeout. The default is 100 seconds. The valid range is 1 to 86,400.
IP Reputation	Enable to apply the FortiGuard IP reputation service. See Managing IP Reputation policy settings .
Customized SSL Ciphers Flag	Enable or disable the Customized SSL Ciphers Flag.
Geo IP Block List	Select a Geo IP block list configuration object. See Using the Geo IP block list .
Geo IP Whitelist	Select a whitelist configuration object. See Using the Geo IP whitelist .
HTTP	
Client Timeout	Client-side TCP connection timeout. The default is 50 seconds. The valid range is 1 to 3,600.
Server Timeout	Server-side IP session timeout. The default is 50 seconds. The valid range is 1 to 3,600.
Connect Timeout	Multiplexed server-side TCP connection timeout. Usually less than the client-side timeout. The default is 5 seconds. The valid range is 1 to 3,600.
Queue Timeout	Specifies how long connection requests to a backend server remain in a queue if the server has reached its maximum number of connections. If the timeout period expires before the client can connect, FortiADC drops the connection and sends a 503 error to the client. The default is 5 seconds. The valid range is 1 to 3,600.

Type	Profile Configuration Guidelines
HTTP Request Timeout	Client-side HTTP request timeout. The default is 50 seconds. The valid range is 1 to 3,600.
HTTP Keepalive Timeout	The default is 50 seconds. The valid range is 1 to 3,600.
Buffer Pool	Enable or disable buffering.
Source Address	Use the original client IP address as the source address when connecting to the real server.
X-Forwarded-For	Append the client IP address found in IP layer packets to the HTTP header that you have specified in the X-Forwarded-For Header setting. If there is no existing X-Forwarded-For header, the system creates it.
X-Forwarded-For Header	Specify the HTTP header to which to write the client IP address. Typically, this is the X-Forwarded-For header, but it is customizable because you might support traffic that uses different headers for this. Do not include the 'X-' prefix. Examples: Forwarded-For, Real-IP, or True-IP.
IP Reputation	Enable to apply the FortiGuard IP reputation service. See Managing IP Reputation policy settings .
HTTP Mode	<ul style="list-style-type: none"> • Server Close—Close the connection to the real server after each HTTP transaction. • Once Only— An HTTP transaction can consist of multiple HTTP requests (separate requests for an HTML page and the images contained therein, for example). To improve performance, the "once only" flag instructs the FortiADC to evaluate only the first set of headers in a connection. Subsequent requests belonging to the connection are not load balanced, but sent to the same server as the first request. • Keep Alive—Do not close the connection to the real server after each HTTP transaction. Instead, keep the connection between FortiADC and the real server open until the client-side connection is closed. This option is required for applications like Microsoft SharePoint.
Customized SSL Ciphers Flag	Enable or disable the Customized SSL Ciphers Flag.
Compression	Select a compression configuration object. See Configuring compression rules .
Caching	Select a caching configuration object. See Using caching features .
Geo IP Block List	Select a Geo IP block list configuration object. See Using the Geo IP block list .

Type	Profile Configuration Guidelines
Geo IP Whitelist	Select a whitelist configuration object. See Using the Geo IP whitelist .
Geo IP Redirect URL	For HTTP, if you have configured a Geo IP redirect action, specify a redirect URL.
FTP	
Timeout TCP Session	Client-side timeout for connections where the client has not sent a FIN signal, but the connection has been idle. The default is 100 seconds. The valid range is 1 to 86,400.
Timeout TCP Session after FIN	Client-side connection timeout. The default is 100 seconds. The valid range is 1 to 86,400.
IP Reputation	Enable to apply the FortiGuard IP reputation service. See Managing IP Reputation policy settings .
Customized SSL Ciphers Flag	Enable or disable the Customized SSL Ciphers Flag.
Geo IP Block List	Select a Geo IP block list configuration object. See Using the Geo IP block list .
Geo IP Whitelist	Select a whitelist configuration object. See Using the Geo IP whitelist .
RADIUS	
Customized SSL Ciphers Flag	Enable or disable the Customized SSL Ciphers Flag.
Timeout RADIUS Session	The default is 300 seconds. The valid range is 1 to 3,600.
Geo IP Block List	Select a Geo IP block list configuration object. See Using the Geo IP block list .
Geo IP Whitelist	Select a whitelist configuration object. See Using the Geo IP whitelist .
RDP	
Client Timeout	Client-side TCP connection timeout. The default is 50 seconds. The valid range is 1 to 3,600.
Server Timeout	Server-side IP session timeout. The default is 50 seconds. The valid range is 1 to 3,600.

Type	Profile Configuration Guidelines
Connect Timeout	Multiplexed server-side TCP connection timeout. Usually less than the client-side timeout. The default is 5 seconds. The valid range is 1 to 3,600.
Queue Timeout	Specifies how long connection requests to a backend server remain in a queue if the server has reached its maximum number of connections. If the timeout period expires before the client can connect, FortiADC drops the connection and sends a 503 error to the client. The default is 5 seconds. The valid range is 1 to 3,600.
Buffer Pool	Enable or disable buffering.
Source Address	Use the original client IP address as the source address in the connection to the real server.
IP Reputation	Enable to apply the FortiGuard IP reputation service. See Managing IP Reputation policy settings .
Customized SSL Ciphers Flag	Enable or disable the Customized SSL Ciphers Flag.
Geo IP Block List	Select a Geo IP block list configuration object. See Using the Geo IP block list .
Geo IP Whitelist	Select a whitelist configuration object. See Using the Geo IP whitelist .
TCPS	
Client Timeout	Client-side TCP connection timeout. The default is 50 seconds. The valid range is 1 to 3,600.
Server Timeout	Server-side IP session timeout. The default is 50 seconds. The valid range is 1 to 3,600.
Connect Timeout	Multiplexed server-side TCP connection timeout. Usually less than the client-side timeout. The default is 5 seconds. The valid range is 1 to 3,600.
Queue Timeout	Specifies how long connection requests to a backend server remain in a queue if the server has reached its maximum number of connections. If the timeout period expires before the client can connect, the system drops the connection and sends a 503 error to the client. The default is 5 seconds. The valid range is 1 to 3,600.
Buffer Pool	Enable or disable buffering.

Type	Profile Configuration Guidelines
Source Address	Use the original client IP address as the source address in the connection to the real server.
IP Reputation	Enable to apply the FortiGuard IP reputation service. See Managing IP Reputation policy settings .
Customized SSL Ciphers Flag	Enable or disable the use of user-specified cipher suites.
Customized SSL Ciphers	<p>If the customize cipher flag is enabled, specify a colon-separated, ordered list of cipher suites.</p> <p>An empty string is allowed. If empty, the default cipher suite list is used.</p>

Type	Profile Configuration Guidelines
SSL Ciphers	<p>Ciphers are listed from strongest to weakest:</p> <ul style="list-style-type: none"> • ECDHE-ECDSA-AES256-GCM-SHA384 • ECDHE-ECDSA-AES256-SHA384 • ECDHE-ECDSA-AES256-SHA • ECDHE-ECDSA-AES128-GCM-SHA256 • ECDHE-ECDSA-AES128-SHA256 • ECDHE-ECDSA-AES128-SHA • ECDHE-ECDSA-DES-CBC3-SHA • ECDHE-ECDSA-RC4-SHA • ECDHE-RSA-AES256-GCM-SHA384 • ECDHE-RSA-AES256-SHA384 • ECDHE-RSA-AES256-SHA • DHE-RSA-AES256-GCM-SHA384 • DHE-RSA-AES256-SHA256 • DHE-RSA-AES256-SHA • AES256-GCM-SHA384 • AES256-SHA256 • AES256-SHA • ECDHE-RSA-AES128-GCM-SHA256 • ECDHE-RSA-AES128-SHA256 • ECDHE-RSA-AES128-SHA • DHE-RSA-AES128-GCM-SHA256 • DHE-RSA-AES128-SHA256 • DHE-RSA-AES128-SHA • AES128-GCM-SHA256 • AES128-SHA256 • AES128-SHA • ECDHE-RSA-RC4-SHA • RC4-SHA • RC4-MD5 • ECDHE-RSA-DES-CBC3-SHA • EDH-RSA-DES-CBC3-SHA • DES-CBC3-SHA • EDH-RSA-DES-CBC-SHA • DES-CBC-SHA • eNULL <p>We recommend retaining the default list. If necessary, you can deselect the SSL ciphers that you do not want to support.</p>

Type	Profile Configuration Guidelines
Allow SSL Versions	<p>You have the following options:</p> <ul style="list-style-type: none"> • SSLv2 • SSLv3 • TLSv1.0 • TLSv1.1 • TLSv1.2 <p>We recommend retaining the default list. If necessary, you can deselect SSL versions you do not want to support.</p> <p>Note: FortiADC does not support session reuse for SSLv2 at the client side. Instead, a new SSL session is started.</p>
Client SNI Required	Require clients to use the TLS server name indication (SNI) extension to include the server hostname in the TLS client hello message. Then, the FortiADC system can select the appropriate local server certificate to present to the client.
Geo IP Block List	Select a Geo IP block list configuration object. See Using the Geo IP block list .
Geo IP Whitelist	Select a whitelist configuration object. See Using the Geo IP whitelist .
Local Certificate Group	A configuration group that includes the certificates this virtual server presents to SSL/TLS clients. This should be the backend servers' certificate, NOT the appliance's GUI web server certificate. See Manage certificates .
Certificate Verify	Select a certificate validation policy. See Manage and validate certificates .
HTTPS	
HTTPS	<p>Same as HTTP, plus the certificate settings listed next.</p> <p>See Chapter 16: SSL Transactions for an overview of HTTPS features.</p>
SSL Proxy Mode	Enable or disable SSL forward proxy.
Customized SSL Ciphers Flag	Enable or disable use of user-specified cipher suites.
Customized SSL Ciphers	<p>If the customize cipher flag is enabled, specify a colon-separated, ordered list of cipher suites.</p> <p>An empty string is allowed. If empty, the default cipher suite list is used.</p>

Type	Profile Configuration Guidelines
SSL Ciphers	We recommend retaining the default list. If necessary, you can deselect ciphers you do not want to support.
Allow SSL Versions	<p>We recommend retaining the default list. If necessary, you can deselect SSL versions you do not want to support.</p> <p>Note: FortiADC does not support session reuse for SSLv2 at the client side. Instead, a new SSL session is started.</p>
Client SNI Required	Require clients to use the TLS server name indication (SNI) extension to include the server hostname in the TLS client hello message. Then, the FortiADC system can select the appropriate local server certificate to present to the client.
Local Certificate Group	A configuration group that includes the certificates this virtual server presents to SSL/TLS clients. This should be the backend servers' certificate, NOT the appliance's GUI web server certificate. See Manage certificates .
Certificate Verify	Select a certificate validation policy. See Manage and validate certificates .
TURBO HTTP	
Timeout TCP Session	Client-side timeout for connections where the client has not sent a FIN signal, but the connection has been idle. The default is 100 seconds. The valid range is 1 to 86,400.
Timeout TCP Session after FIN	Client-side connection timeout. The default is 100 seconds. The valid range is from 1 to 86,400.
IP Reputation	Enable to apply the FortiGuard IP reputation service.
Customized SSL Ciphers Flag	Enable or disable the Customized SSL Ciphers Flag.
Geo IP Block List	Select a Geo IP block list configuration object. See Using the Geo IP block list .
Geo IP Whitelist	Select a whitelist configuration object. See Using the Geo IP whitelist .
SIP	
SIP Max Size	Maximum message size. The default is 65535 bytes. The valid range is from 1 to 65,535.

Type	Profile Configuration Guidelines
Server Keepalive Timeout	Maximum wait for a new server-side request to appear. The default is 30 seconds. The valid range is 5-300.
Server Keepalive	Enable/disable a keepalive period for new server-side requests. Supports CRLF ping-pong for TCP connections. Enabled by default.
Client Keepalive	Enable/disable a keepalive period for new client-side requests. Supports CRLF ping-pong for TCP connections. Disabled by default.
Client Protocol	Client-side transport protocol: <ul style="list-style-type: none"> • TCP • UDP (default)
Server Protocol	Server-side transport protocol. <ul style="list-style-type: none"> • TCP • UDP Default is "unset", so the client-side protocol determines the server-side protocol.
Failed Client Type	Action when the SIP client cannot be reached: <ul style="list-style-type: none"> • Drop—Drop the connection. • Send—Drop the connection and send a message, for example, a status code and error message.
Failed Server Type	Action when the SIP server cannot be reached: <ul style="list-style-type: none"> • Drop—Drop the connection. • Send—Drop the connection and send a message, for example, a status code and error message.
Insert Client IP	Enable/disable option to insert the client source IP address into the X-Forwarded-For header of the SIP request.
Client-Request-Header-Insert (maximum 4 members)	
Type	<ul style="list-style-type: none"> • Insert If Not Exist—Insert before the first header only if the header is not already present. • Insert Always—Insert before the first header even if the header is already present. • Append If Not Exist—Append only if the header is not present. • Append Always—Append after the last header.
HeaderName:Value	The header:value pair to be inserted.

Type	Profile Configuration Guidelines
Client-Request-Header-Erase (maximum 4 members)	
Type	<ul style="list-style-type: none"> • All—Parse all headers for a match. • First—Parse the first header for a match.
HeaderName	Header to be erased.
Client-Response-Header-Insert (maximum 4 members)	
Type	<ul style="list-style-type: none"> • Insert If Not Exist—Insert before the first header only if the header is not already present. • Insert Always—Insert before the first header even if the header is already present. • Append If Not Exist—Append only if the header is not present. • Append Always—Append after the last header.
HeaderName:Value	The header:value pair to be inserted.
Client-Response-Header-Erase (maximum 4 members)	
Type	<ul style="list-style-type: none"> • All—Parse all headers for a match. • First—Parse the first header for a match.
HeaderName	Header to be erased.
Server-Request-Header-Insert (maximum 4 members)	
Type	<ul style="list-style-type: none"> • Insert If Not Exist—Insert before the first header only if the header is not already present. • Insert Always—Insert before the first header even if the header is already present. • Append If Not Exist—Append only if the header is not present. • Append Always—Append after the last header.
HeaderName:Value	The header:value pair to be inserted.
Server-Request-Header-Erase (maximum 4 members)	
Type	<ul style="list-style-type: none"> • All—Parse all headers for a match. • First—Parse the first header for a match.
HeaderName	Header to be erased.
Server-Response-Header-Insert (maximum 4 members)	

Type	Profile Configuration Guidelines
Type	<ul style="list-style-type: none"> • Insert If Not Exist—Insert before the first header only if the header is not already present. • Insert Always—Insert before the first header even if the header is already present. • Append If Not Exist—Append only if the header is not present. • Append Always—Append after the last header.
HeaderName:Value	The header:value pair to be inserted.
Server-Response-Header-Erase (maximum 4 members)	
Type	<ul style="list-style-type: none"> • All—Parse all headers for a match. • First—Parse the first header for a match.
HeaderName	Header to be erased.
SMTP	
Starttls Active Mode	<p>Select one of the following:</p> <ul style="list-style-type: none"> • Allow—The client can either use or not use the STARTTLS command. • Require—The STARTTLS command must be used to encrypt the connection first. • None—The STARTTLS command is NOT supported.
Forbidden Command	<p>Select any, all, or none of the commands (i.e., expn, turn, vrfy).</p> <p>If selected, the command or commands will be rejected by FortiADC; otherwise, the command or commands will be accepted and forwarded to the back end.</p>
Domain Name	Specify the domain name.
Local Certificate Group	LOCAL_CERT_GROUP.
Certificate Verify	Specify the certificate verify configuration object.
RTMP	
Source Address	When enabled, specify the client address to be used to connect to the server pool.
RTSP	
Max Header Size	Specify the maximum size of the RTSP header.

Type	Profile Configuration Guidelines
Source Address	When enabled, specify the client address to be used to connect to the server pool.
MySQL	Note: The system does not provide default MySQL profiles as it does with the other protocols.
Single Master	If selected, the profile will use the single-master mode. You will then need to specify and configure the master server and slave servers.
Sharding	If selected, the profile will use the sharding mode to load-balance MySQL traffic.

Configuring error pages

When backend real servers are unavailable, the system can respond to clients attempting HTTP/HTTPS connections with either an error message or an HTML error page.

You do not have to create an error message configuration object if you want only to send an error message. You can configure an error message from within the virtual server configuration.

After you have created an error page configuration object, you can select it in the virtual server configuration.

Before you begin:

Copy the error message file to a location you can reach from your browser. The error message file must be named `index.html` and must be contained in a zip file.

- You must have Read-Write permission for Load Balance settings.

To upload an error message file:

1. Go to Server Load Balance > Application Resources.
2. Click the **Error Page** tab.
3. Click **Add** to display the configuration editor.
4. Give the configuration object a name, browse and select the error message zip file, and click the upload icon.
5. Save the configuration.

Using source pools

This topic includes a procedure for configuring the source IP address pools used in NAT, and examples of NAT deployments. It includes the following sections:

- [Configuring source pools](#)
- [Example: DNAT](#)
- [Example: full NAT](#)
- [Example: NAT46 \(Layer 4 virtual servers\)](#)

- [Example: NAT64 \(Layer 4 virtual servers\)](#)
- [Example: NAT46 \(Layer 7 virtual servers\)](#)
- [Example: NAT64 \(Layer 7 virtual servers\)](#)

Configuring source pools

You use the Source Pool page to create configuration objects for source IP addresses used for NAT in Layer 4 virtual server configurations.

In a Layer 4 virtual server configuration, you select a “packet forwarding method” that includes the following network address translation (NAT) options:

- Direct Routing—Does not rewrite source or destination IP addresses.
- DNAT—Rewrites the destination IP address for packets before it forwards them.
- Full NAT—Rewrites both the source and destination IP addresses. Use for standard NAT, when client and server IP addresses are all IPv4 or all IPv6.
- NAT46—Rewrites both the source and destination IP addresses. Use for NAT 46, when client IP addresses are IPv4 and server IP addresses are IPv6.
- NAT64—Rewrites both the source and destination IP addresses. Use for NAT 64, when client IP addresses are IPv6 and server IP addresses are IPv4.

In a Layer 7 virtual server configuration, you do not select a packet forwarding option. Layer 7 virtual servers use NAT46 and NAT64 to support those traffic flows, but they do not use the Source Pool configuration.

See the examples that follow the procedure for illustrated usage.

Before you begin:

- You must have a good understanding of NAT. You must know the address ranges your network has provisioned for NAT.
- Be sure to configure the backend servers to use the FortiADC address as the default gateway so that server responses are also rewritten by the NAT module.
- You must have Read-Write permission for Load Balance settings.

After you have configured a source pool IP address range configuration object, you can select it in the virtual server configuration. You can assign a virtual server multiple source pools (with the same or different source pool interface associated with it).

To configure a source pool:

1. Go to Server Load Balance > Virtual Server.
2. Click the **NAT Source Pool** tab.
3. Click **Add** to display the configuration editor.
4. Complete the configuration as described in [Table 21](#).
5. Save the configuration.

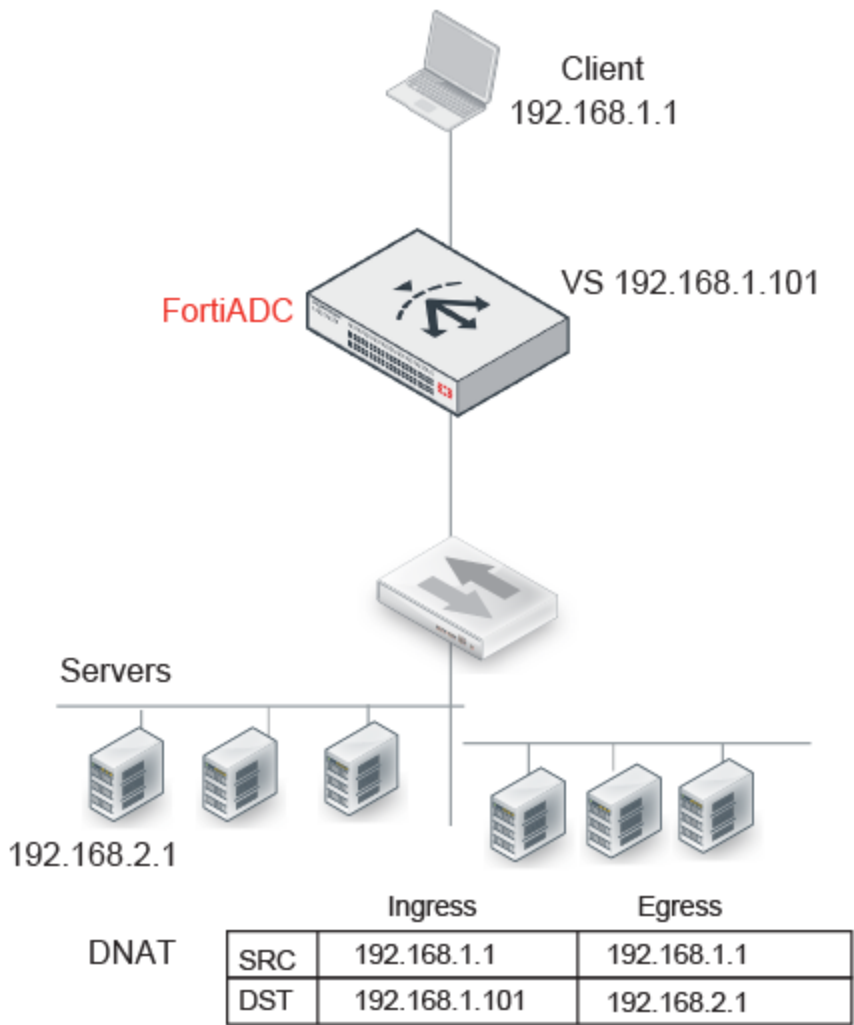
Table 21: Source pool configuration

Settings	Guidelines
Name	<p>Configuration name. Valid characters are A-Z, a-z, 0-9, _, and -. No spaces. You reference this name in the virtual server configuration.</p> <p>Note: After you initially save the configuration, you cannot edit the name.</p>
Interface	Interface to receive responses from the backend server. The interface used for the initial client traffic is determined by the virtual server configuration.
Address Type	<ul style="list-style-type: none"> • IPv4 • IPv6
Address Range	The first address in the address pool.
To	The last address in the address pool.
Node Member	
Name	<p>Create a node member list to be used in an HA active-active deployment. In an active-active deployment, node interfaces are configured with a list of IP addresses for all nodes in the cluster. You use this configuration to provision SNAT addresses for each of the nodes.</p> <p>Name is a configuration name. Valid characters are A-Z, a-z, 0-9, _, and -. No spaces. You reference this name in the virtual server configuration.</p> <p>Note: After you initially save the configuration, you cannot edit the name.</p>
Pool Type	IPv4 or IPv6.
Minimum IP	The first address in the address pool.
Maximum IP	The last address in the address pool.
Interface	Interface to receive responses from the backend server. The interface used for the initial client traffic is determined by the virtual server configuration.
HA Node Number	Specify the HA cluster node ID.

Example: DNAT

Figure 31 illustrates destination NAT (DNAT). The NAT module rewrites only the destination IP address. Therefore, if you configure destination NAT, you do not need to configure a source pool. In this DNAT example, the destination IP address in the packets it receives from the client request is the IP address of the virtual server—192.168.1.101. The NAT module translates this address to the address of the real server selected by the load balancer—in this example, 192.168.2.1. The system maintains this NAT table and performs the inverse translation when it receives the server-to-client traffic.

Figure 31: Destination NAT




Example: full NAT

Figure 32 illustrates full NAT. The source IP / destination IP pair in the packets received is SRC 192.168.1.1 / DST 192.168.1.101. The NAT module translates the source IP address to the next available address in the source pool—in this example, 192.168.2.101. It translates the destination IP address to the address of the real server selected by the load balancer—in this example, 192.168.2.1.

The system maintains this NAT table and performs the inverse translation when it receives the server-to-client traffic.

Figure 32: Full NAT



Example: NAT46 (Layer 4 virtual servers)

Figure 33 illustrates full NAT with NAT46. The IPv6 client connects to the virtual server IPv4 address. The source IP / destination IP pair in the packets received is SRC 192.168.1.1 / DST 192.168.1.101. The NAT module translates the source IP address to the next available IPv6 address in the source pool—in this example, 2002::2:1001. It translates the destination IP address to the IPv6 address of the real server selected by the load balancer—in this example, 2002::2:1.

The system maintains this NAT table and performs the inverse translation when it receives the server-to-client traffic.

Figure 33: NAT46 (Layer 4 virtual servers)

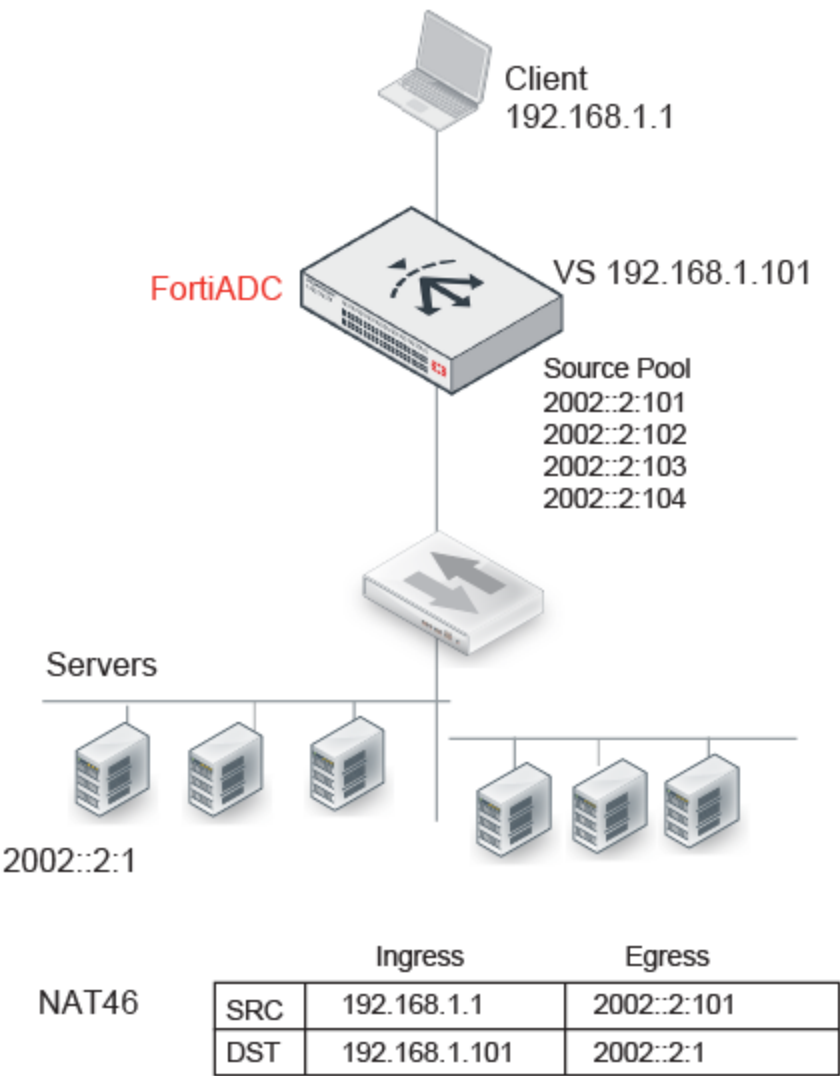


Table 22: Limitations: NAT46 (Layer 4 virtual servers)

Features	Notes
Profile	Not Supported: FTP
ICMP	ICMP traffic is dropped.

Example: NAT64 (Layer 4 virtual servers)

Figure 34 illustrates full NAT with NAT64. The IPv6 client connects to the virtual server IPv6 address. The source IP / destination IP pair in the packets received is SRC 2001::1:1 / DST 2001::1:101. The NAT module translates the source IP address to the next available IPv4 address in the source pool—in this example, 192.168.2.101. It translates the destination IP address to the IPv4 address of the real server selected by the load balancer—in this example, 192.168.2.1.

The system maintains this NAT table and performs the inverse translation when it receives the server-to-client traffic.

Figure 34: NAT64 (Layer 4 virtual servers)

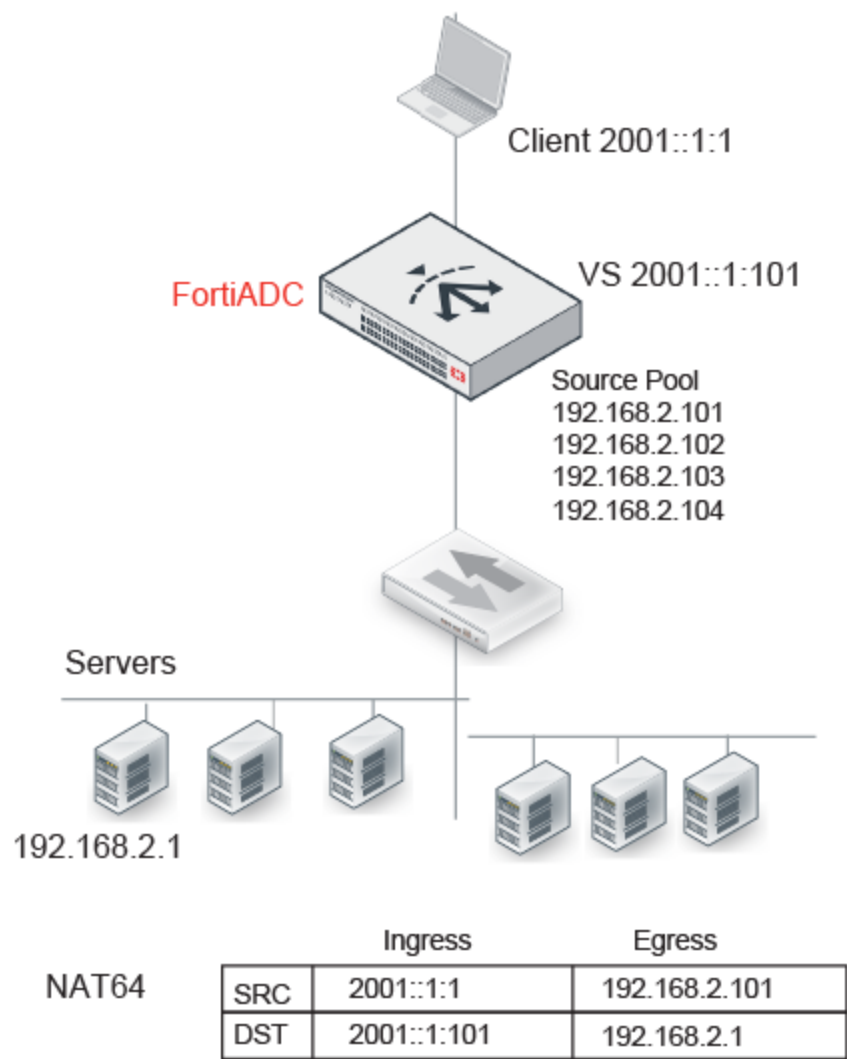


Table 23: Limitations: NAT64 (Layer 4 virtual servers)

Features	Notes
Profiles	Not Supported: FTP
ICMP	ICMP traffic is dropped.
Security	Not Supported: IP Reputation, DoS protection, Security logs and reports

Example: NAT46 (Layer 7 virtual servers)

Figure 35 illustrates full NAT with NAT46. The IPv4 client connects to the virtual server IPv4 address. The source IP / destination IP pair in the packets received is SRC 192.168.1.1 / DST 192.168.1.101. The NAT module translates the source IP address to the IPv6 address of the egress interface that has IPv6 connectivity with the real server—in this example, 2002::2:1001. It translates the destination IP address to the IPv6 address of the real server selected by the load balancer—in this example, 2002::2:1.

The system maintains this NAT table and performs the inverse translation when it receives the server-to-client traffic.

Figure 35: NAT46 (Layer 7 virtual servers)




Table 24: Limitations: NAT46 (Layer 7 virtual servers)

Feature	Note
Profiles	Not Supported: RADIUS, HTTP Turbo
Profile options	Not supported: Source Address (Using the original source IP address for the connection to the real server is contrary to the purpose of NAT.)
Virtual server options	Not supported: Connection Rate Limit
Real server pool options	Not supported: Connection Rate Limit

Example: NAT64 (Layer 7 virtual servers)

Figure 36 illustrates full NAT with NAT64. The IPv6 client connects to the virtual server IPv6 address. The source IP / destination IP pair in the packets received is SRC 2001::1:1 / DST 2001::1:101. The NAT module translates the source IP address to the IPv4 address of the egress interface that has IPv4 connectivity with the real server—in this example, 192.168.2.101. It translates the destination IP address to the IPv4 address of the real server selected by the load balancer—in this example, 192.168.2.1.

The system maintains this NAT table and performs the inverse translation when it receives the server-to-client traffic.

Figure 36: NAT64 (Layer 7 virtual servers)

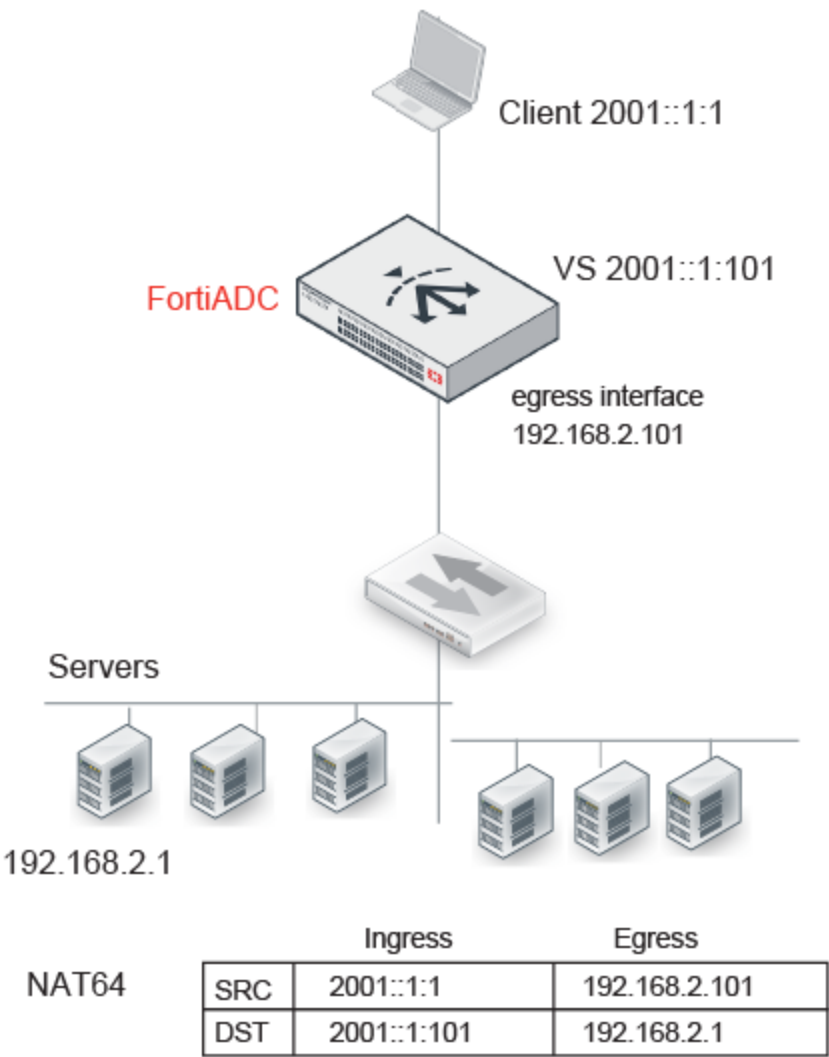


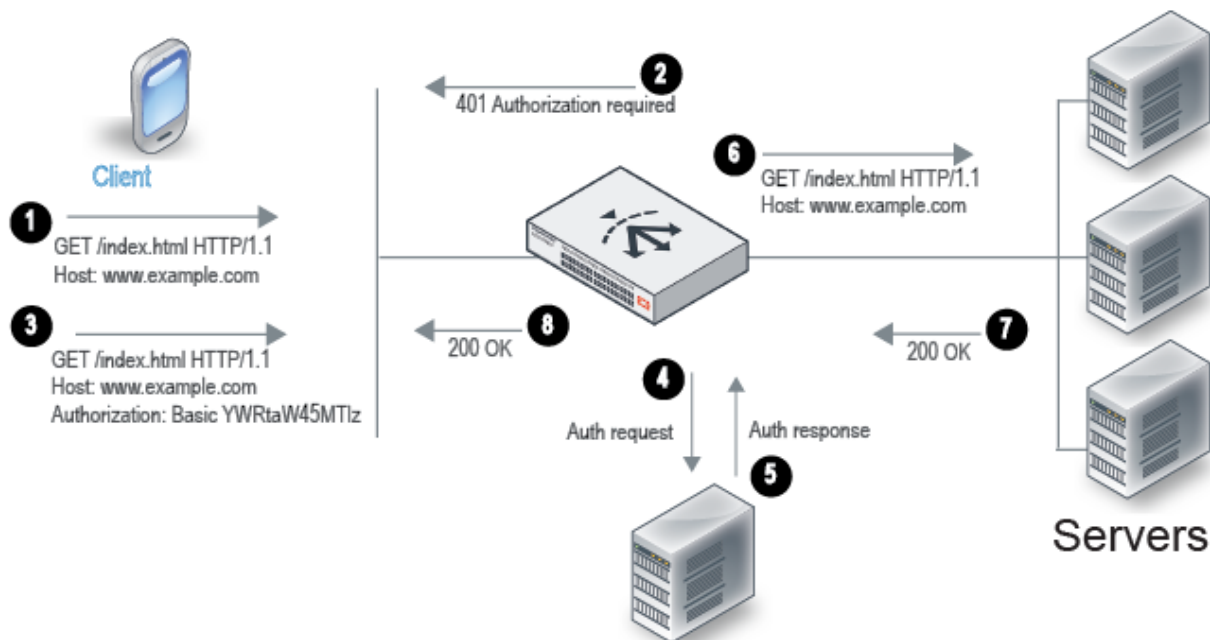
Table 25: Limitations: NAT64 (Layer 7 virtual servers)

Feature	Note
Profiles	Not Supported: RADIUS, HTTP Turbo
Profile options	Not supported: Source Address (Using the original source IP address for the connection to the real server is contrary to the purpose of NAT.)
Virtual server options	Not supported: Connection Rate Limit
Real server pool options	Not supported: Connection Rate Limit
Security	Not Supported: IP Reputation, DoS protection, Security logs and reports

Configuring authentication policies

Auth policies set the conditions that mandate authentication and reference the user group that has authorization. For example, you can define an auth policy that has the following logic: if the Host header matches example.com and the URI matches /index.html, then the group example-group is authorized. FortiADC supports the Basic Authentication Scheme described in [RFC 2617](#).

Figure 37 illustrates the client-server communication when authorization is required.

Figure 37: Authorization and authentication

1. The client sends an HTTP request for a URL belonging to a FortiADC virtual server that has an authorization policy.
2. FortiADC replies with an HTTP 401 to require authorization. On the client computer, the user might be prompted with a dialog box to provide credentials.

3. The client reply includes an [Authorization](#) header that gives the credentials.
4. FortiADC sends a request to the server (local, LDAP, or RADIUS) to authenticate the user.
5. The authentication server sends its response, which can be cached according to your user group configuration.
6. If authentication is successful, FortiADC continues processing the traffic and forwards the request to the real server.
7. The real server responds with an HTTP 200 OK.
8. FortiADC processes the traffic and forwards the server response to the client.

Before you begin:

- You must have created the user groups to be authorized with the policy. You also configure users and authentication servers separately. See [Configuring user groups](#).
- You must have read-write permission for Server Load Balance settings.

After you have configured an auth policy, you can select it in the virtual server configuration. Note the following requirements:

- Virtual server type must be Layer 2 or Layer 7.
- Profile type must be HTTP or HTTPS.
- The profile option once-only must be disabled.

To configure an authentication policy:

1. Go to Server Load Balance > Application Resources.
2. Click the **Authentication Policy** tab.
3. Click **Add** to display the configuration editor.
4. Complete the configuration as described in [Table 26](#).
5. Save the configuration.

Table 26: Authentication policy configuration

Settings	Guidelines
Name	Configuration name. Valid characters are A-Z, a-z, 0-9, _, and -. No spaces. You reference this name in the virtual server configuration. Note: After you initially save the configuration, you cannot edit the name.
Member	
Host Status	If enabled, require authorization only for the specified host. If disabled, ignore hostname in the HTTP request header and require authorization for requests with any Host header. Disabled by default.
Host	If Host Status is enabled, the policy matches only if the Host header matches this value. Complete, exact matching is required. For example, <code>www.example.com</code> matches <code>www.example.com</code> but not <code>www.example.com.hk</code> . Note: This field becomes available only if Host Status is enabled.

Settings	Guidelines
Type	Select either of the following: <ul style="list-style-type: none"> Standard SAML
User Realm	Realm to which the Path URI belongs. The realm is included in the basic authentication header in the HTTP 401 message sent to the client. If a request is authenticated and a realm specified, the same credentials are deemed valid for other requests within this realm.
Path	Require authorization only if the URI of the HTTP request matches this pathname. If none is specified, requests to any URI require authorization. The value is parsed as a match string prefix. For example, <code>/abc</code> matches <code>http://www.example.com/abcd</code> and <code>http://www.example.com/abc/11.html</code> but not <code>http://www.example.com/1abcd</code> .
User Group	Select the user group that is authorized to access the protected resource.
SAML SSO ID	Select a SAML SSO ID. Note: This field becomes available only when authentication policy type is set to SAML. See Type above.

Configuring load-balancing (LB) methods

The system includes predefined configuration objects for all supported load balancing methods, and there is no need to create additional configuration objects. You may choose to do so, however, for various reasons, for example, to use a naming convention that makes the purpose of the configuration clear to other administrators.

Table 27 describes the predefined methods.

Table 27: Predefined LB methods

Predefined	Description
LB_METHOD_ROUND_ROBIN	Client requests are distributed across a pool of servers sequentially.
LB_METHOD_LEAST_CONNECTION	New client requests are sent to servers with the fewest current connections to clients.
LB_METHOD_FASTEST_RESPONSE	New requests are forwarded to the server with the fastest response to health check tests.

Predefined	Description
LB_METHOD_URI	Selects the server based on a hash of the URI found in the HTTP header, excluding hostname.
LB_METHOD_FULL_URI	Selects the server based on a hash of the full URI string found in the HTTP header. The full URI string includes the hostname and path.
LB_METHOD_HOST	Selects the server based on a hash of the hostname in the HTTP Request header Host field.
LB_METHOD_HOST_DOMAIN	Selects the server based on a hash of the domain name in the HTTP Request header Host field.
LB_METHOD_DEST_IP_HASH	Selects the next hop based on a hash of the destination IP address. This method can be used with Layer-2 virtual servers.
LB_METHOD_DYNAMIC_LOAD	<p>Selects the server with the highest weight assigned to it based on its SNMP health check.</p> <p>Note: Dynamic load-balancing is a load-balancing method in which FortiADC (the load-balancer) actively polls server pool members, and then assigns a weighted value to each member based on a set of default or user-defined thresholds. The value ranges from 1 to 256, and determines the amount of traffic FortiADC directs to a member. The greater the value that FortiADC assigns to a member, the more client requests it (the member) receives.</p> <p>Dynamic load-balancing relies on the status of SNMP health check to calculate the load on each real server. The health check covers a real server's CPU, memory, and disk usage. When a real server has exceeded its health check thresholds, it will be marked as "down". If that happens, FortiADC will stop sending client requests to that server.</p>

Before you begin:

- You must have read-write permission for load balance settings.

To configure a load-balancing method configuration object:

1. Go to Server Load Balance > Application Resources.
2. Click the **LB Method** tab.
3. Click **Add** to display the configuration editor.
4. Specify a unique a name for the configuration object.
5. Select a load-balancing method from the list menu.
6. Save the configuration.

Configuring an L2 exception list

In some jurisdictions, SSL interception and decryption is disfavored for some types of websites or disallowed entirely. You use the L2 Exception List configuration to define such destinations. You can leverage FortiGuard web filter categories, and you can configure a list of additional destinations.

Before you begin:

- You must have created a Web Filter Profile configuration that includes the web categories to exclude from SSL decryption.
- You must have hostname or IP address details on additional destinations you want to exclude from SSL decryption.
- You must have Read-Write permission for Load Balance settings.

After you have created an L2 exception list configuration object, you can select it in a Layer 2 virtual server configuration.

To configure an exception list:

1. Go to Server Load Balance > SSL-FP Resources.
2. Click the **L2 Exception List** tab.
3. Click **Add** to display the configuration editor.
4. Complete the configuration as described in [Table 28](#).
5. Save the configuration.

Table 28: L2 exception list configuration

Settings	Guidelines
Name	Configuration name. Valid characters are A-Z, a-z, 0-9, _, and -. No spaces. You reference this name in the profile configuration. Note: After you initially save the configuration, you cannot edit the name.
Description	A string to describe the purpose of the configuration, to help you and other administrators more easily identify its use.
Web Filter Profile	Select a Web Filter Profile configuration.
Member	
Type	How you want to define the exception: <ul style="list-style-type: none">• Host• IP
Host Pattern	Specify a wildcard pattern, such as *.example.com.

Settings	Guidelines
IP/Netmask	Specify the IP address and CIDR-formatted subnet mask, separated by a forward slash, such as 192.0.2.0/24.
	Note: <ul style="list-style-type: none"> Dotted quad formatted subnet masks are not accepted. IPv6 addresses are not supported.

Creating a Web Filter Profile configuration

You use the web filter profile configuration to create groups of FortiGuard categories that you want to include in the SSL forward proxy "L2 Exception List" configuration. The web filter profile should include categories that should not be processed by the outbound L2 SSL forward proxy feature. To address privacy concerns, you can include categories such as "Personal Privacy", "Finance and Banking", "Health and Wellness", and Medicine.

Before you begin:

- Learn about FortiGuard web filter categories. Go to <http://fortiguard.com/webfilter>.
- You must have Read-Write permission for Load Balance settings.

After you have created a web filter profile configuration object, you can select it in a L2 exception list configuration.

To create a web filter profile configuration:

- Go to Server Load Balance > SSL-FP Resources.
- Click the **Web Filter Profile** tab.
- Click **Add** to display the configuration editor.
- Complete the configuration as described in [Table 29](#).
- Save the configuration.

Table 29: Web Filter Profile configuration

Settings	Guidelines
Name	Configuration name. Valid characters are A-Z, a-z, 0-9, _, and -. No spaces. You reference this name in the profile configuration.
	Note: After you initially save the configuration, you cannot edit the name.
Description	A string to describe the purpose of the configuration, to help you and other administrators more easily identify its use.
Category-Members	
Category	Select a category or subcategory from the predefined list.

Using the Web Category tab

The Web Category tab displays the web filter categories imported from FortiGuard. You specify web categories when you create web filter groups.

For information on FortiGuard web categories, go to the FortiGuard website:

<http://fortiguard.com/webfilter>

Before you begin:

- You must have read permission for load balancing settings.

To display web categories:

1. Go to Server Load Balance > SSL-FP Resources.
2. Click the **Web Category** tab.

To manage how long the URL lists from FortiGuard are cached:

1. Go to System > FortiGuard.
2. Under Web Filter Configure, adjust caching settings as desired.

Configuring virtual servers

The virtual server configuration supports three classes of application delivery control:

- Layer 7—Persistence, load-balancing, and routing are based on Layer-7 objects, such as HTTP headers, cookies, and so on.
- Layer 4—Persistence, load-balancing, and network address translation are based on Layer-4 objects, such as source and destination IP address.
- Layer 2—This feature is useful when the request's destination IP is unknown and you need to load-balance connections between multiple next-hop gateways.

Before you begin, you must:

- Have a deep understanding of the backend servers and your load-balancing objectives.
- Have configured a real server pool (required) and other configuration objects that you can incorporate into the virtual server configuration, such as persistence rules, user-defined profiles, content routes and rewriting rules, error messages, authentication policies, and source IP address pools if you are deploying NAT.
- Have read-write permission for load-balance configurations.



Unlike virtual IPs on FortiGate or virtual servers on FortiWeb, virtual servers on FortiADC are activated as soon as they are configured and their status are set to **Enable**. You do not need to apply them by selecting them in a policy.

Two Options for virtual server configuration

FortiADC provides two options for configuring virtual servers—Basic Mode and Advanced Mode.

In Basic Mode, you are required to specify only the basic parameters needed to configure a virtual server. FortiADC automatically configures those advanced parameters using the default values when you click the **Save** button. The Basic Mode is for less experienced users who may not have the skills required to configure the advanced features on their own.

The Advanced Mode, on the other hand, is ideal for experienced or "power" users who are knowledgeable and comfortable enough to configure all the advanced features, in addition to the basic ones, on their own.

All virtual servers you have added, whether they are configured through the Basic Mode or the Advanced Mode, end up on the Load Balance > Virtual Server page. You can view the configuration details of a virtual server by clicking the entry.

Basic virtual server configuration

This option is used mostly for beginners who have less experience with FortiADC.

To configure a virtual server using the Basic Mode:

1. From the menu bar, click Server Load Balance >Virtual Server.
2. Click **Add >Basic Mode** to open the Basic Mode configuration editor.
3. Complete the configuration as described in Table 30.
4. Click **Save**.

Table 30: Virtual server configuration Basic Mode

Settings	Guidelines
Name	<p>Specify the virtual server name. Valid characters are A-Z, a-z, 0-9, _, and -. No space is allowed. This name appears in reports and in logs as the SLB "policy".</p> <p>Note: Once saved, the name of a virtual server cannot be changed.</p>

Settings	Guidelines
Application	<p>Select one of the following applications:</p> <ul style="list-style-type: none"> • Microsoft SharePoint Application • Microsoft Exchange Server Application • IIS • Apache • Windows Remote Desktop • HTTPS • TCPS • HTTP Turbo • RADIUS • DNS • SIP • TCP • UDP • FTP • IP • RTSP • RTMP • SMTP
Address	Specify the IP address provisioned for the virtual server.
Port	<p>Accept the default port or specify a port , ports, or a range of ports of your preference.</p> <p>Note: The virtual server will use the specified port or ports to listen for client requests. You can specify up to eight ports or port ranges separated by space. Valid values are from 0 to 65535. Port 0 applies to Layer-4 virtual servers only,</p>
Interface	Select a network interface from the list menu, or specify a new one.
Real Server Pool	Select a real server pool (if you have one already configured) or create a new one.
SSL	This field becomes available as an option only when Application is set to HTTPS. It is disabled by default. You can select the box to enable SSL, which will then requires you to select a Client SSL Profile. See below.
Client SSL Profile	<p>Note: This setting applies to HTTPS, TCPS, and SMTP applications only. In the case of HTTPS, it becomes available only when SSL is enabled.</p> <p>Select a client SSL profile from the drop-down menu.</p>

Settings	Guidelines
Protocol	Note: This setting becomes available only when Application is set to IP . Enter up to eight numeric values or value ranges corresponding to the protocols you'd like to use, separated by space.
Domain Name	Note: This field becomes available only when Application is set to SMTP . Specify the FQDN.

Advanced virtual server configuration

This option is used mostly by advanced users of FortiADC.

To configure a virtual server using the Advanced Mode:

1. From the menu bar, click Server Load Balance > Virtual Server.
2. Click **Add > Advanced Mode**.
3. Complete the configuration as described in [Table 30](#).
4. Save the configuration.

Table 31: Virtual server configuration in Advanced Mode

Settings	Guidelines
Virtual Server	
Name	Configuration name. Valid characters are A-Z, a-z, 0-9, _, and -. No space is allowed. Note: Once saved, the name of a virtual server cannot be changed.
Status	<ul style="list-style-type: none"> • Enable—The server can receive new sessions. • Disable—The server does not receive new sessions and closes any current sessions as soon as possible. • Maintain—The server does not receive new sessions but maintains its current connections.
Type	<ul style="list-style-type: none"> • Layer 7—Persistence, load-balancing, and routing are based on Layer-7 objects, such as HTTP headers, cookies, and so on. • Layer 4—Persistence, load-balancing, and network address translation are based on Layer-4 objects, such as source and destination IP addresses. • Layer 2—This feature is useful when the request's destination IP is unknown and you need to load-balance connections between multiple next-hop gateways.

Settings	Guidelines
Address Type	<ul style="list-style-type: none"> • IPv4 • IPv6 <p>Note: This field applies to Layer-7 and Layer-4 virtual servers only. Keep in mind that IPv6 is not supported for FTP, HTTP Turbo, or SIP profile.</p>
Traffic Group	<p>Select a pre-configured traffic group or create a new one.</p> <p>Note: FortiADC will use the default traffic group if you do not specify one of your own.</p>
Specifics	<p>Note: Some of the settings in this part of the GUI apply to both Layer-7 and Layer-4 virtual servers, and some apply to Layer-7 or Layer-4 virtual servers only, but none of them applies to Layer-2 virtual servers.</p>
Content Routing	<p>Enable to route packets to backend servers based on IP address (Layer-4) or HTTP headers (Layer-7 content); select content route configuration objects and put them in order.</p> <p>Note: Content routing rules override static or policy routes.</p>
Content Routing List	<p>Available only when Content Routing is enabled. In that case, select the options from the list of Available Items, or create new ones.</p> <p>Note: You can select multiple content routing rules in the virtual server configuration. Rules that you add are consulted from top to bottom. The first rule to match is applied. If the traffic does not match any of the content routing rule conditions specified in the virtual server configuration, the system behaves unexpectedly. Therefore, it is important that you create a “catch all” rule that has no match conditions. In the virtual server configuration, this rule should be ordered last so it can be used to forward traffic to a default pool.</p> <p>See Configuring content routes.</p> <p>Note: Not supported for SIP profiles.</p>
Content Rewriting	<p>Enable to rewrite HTTP headers; select content rewriting rules and put them in order.</p> <p>Note: This option applies to Layer-7 virtual servers only.</p>

Settings	Guidelines
Content Rewriting List	<p>Available only when Content Rewriting is enabled. In that case, select the options from the list of Available Items, or create new ones.</p> <p>Note: You can select multiple content rewriting rules in the virtual server configuration. Rules that you add are consulted from top to bottom. The first rule to match is applied. If the traffic does not match any of the content rewriting rule conditions, the header is not rewritten.</p> <p>See Using content rewriting rules.</p> <p>Note: Not supported for SIP profiles.</p>
Transaction Rate Limit	<p>Note: This setting applies to Layer-7 virtual servers only. It is not supported for HTTP Turbo profiles.</p> <p>Specify a limit to the number of HTTP requests per second that virtual server can process. The default is 0 (disabled). The valid range is 1 to 1,048,567 transactions per second.</p> <p>The system counts each client HTTP request against the limit. When the HTTP request rate exceeds the limit, the virtual server will send an HTTP 503 error response to the client.</p>
Packet Forwarding Method	<p>Note: This option applies to Layer-4 virtual servers only.</p> <p>Select one of the following packet forwarding methods:</p> <ul style="list-style-type: none"> • Direct Routing—Forwards the source and destination IP addresses with no changes. <ul style="list-style-type: none"> Note: For FTP profiles, when Direct Routing is selected, you must also configure a persistence method. • DNAT—Replaces the destination IP address with the IP address of the backend server selected by the load balancer. <p>The destination IP address of the initial request is the IP address of the virtual server. Be sure to configure FortiADC as the default gateway on the backend server so that the reply goes through FortiADC and can also be translated.</p> • Full NAT—Replaces both the destination and source IP addresses. IPv4 to IPv4 or IPv6 to IPv6 translation. • NAT46—Replaces both the destination and source IP addresses, translating IPv4 addresses to IPv6 addresses. • NAT64—Replaces both the destination and source IP addresses, translating IPv6 addresses to IPv4 addresses. <p>For Full NAT, NAT46, and NAT64, the source IP address is replaced by an IP address from the pool you specify. The destination IP address is replaced with the IP address of the backend server selected by the load-balancer.</p>

Settings	Guidelines
NAT Source Pool List	<p>Note: This setting applies to Layer-4 virtual server configurations when either Full NAT, NAT46, or NAT64 is selected.</p> <p>Select one or more NAT source pool configuration objects. See Using source pools.</p>
General	
Configuration	
Address	<p>Enter an IPv4 or IPv6 address for the virtual server.</p> <p>Note: You do not specify an IP address for a Layer 2 virtual server. A Layer 2 virtual server is not aware of IP addresses. Instead of routing data for a specific destination, this type of server simply forwards data from the specified network interface and port.</p>
Port	<p>Accept the default port or specify a port, ports, or port ranges of your preference.</p> <p>Note: The virtual server will use the specified port or ports to listen for client requests. You can specify up to eight ports or port ranges separated by space. Valid values are from 0 to 65535. Port 0 applies to Layer-4 virtual servers only.</p> <p>The port range option is useful in deployments where it is desirable to have a virtual IP address with a large number of virtual ports, such as data centers or web hosting companies that use port number to identify their specific customers.</p> <p>Statistics and configurations are applied to the virtual port range as a whole and not to the individual ports within the specified port range.</p> <p>Note: If a Layer 2 virtual server is assigned a network interface that uses port 80 or 443, ensure that the HTTPS and HTTP administrative access options are not enabled for the interface. Setting a port range is not supported for FTP, HTTP Turbo, RADIUS, or Layer 2 TCP profiles.</p>
Connection Limit	<p>Limit the number of concurrent connections. The default is 0 (disabled). The valid range is 1 to 1,048,576 concurrent connections.</p> <p>You can apply a connection limit per real server and per virtual server. Both limits are enforced. Attempted connections that are dropped by security rules are not counted.</p> <p>Note: Not supported for FTP or SIP profiles.</p>


Settings	Guidelines
Connection Rate Limit	<p>With Layer 4 profiles, and with the Layer 2 TCP profile, you can limit the number of new connections per second. The default is 0 (disabled). The valid range is 1 to 86,400 connections per second.</p> <p>You can apply a connection rate limit per real server and per virtual server. Both limits are enforced. Attempted connections that are dropped by security rules are not counted.</p> <p>Note: Not supported for FTP profiles.</p>
Interface	Network interface that receives client traffic for this virtual server.
Resources	
Profile	Select a predefined or user-defined profile configuration object. See Configuring Application profiles .
Persistence	Select a predefined or user-defined persistence configuration object. See Configuring persistence rules .
Method	Select a predefined or user-defined method configuration object. See .
Real Server Pool	Select a real server pool configuration object. See Configuring real server pools .
Auth Policy	Select an auth policy configuration object. HTTP/HTTPS only. See Configuring authentication policies .
Scripting	Select a scripting configuration object. HTTP/HTTPS only. See Using scripts .
L2 Exception List	Select an exception configuration object. Layer 2 HTTPS/TCP only. See Configuring an L2 exception list .
Error Page	
Error Page	<p>Select an error page configuration object. See Configuring error pages.</p> <p>Note: This setting is not supported for SIP profiles.</p>
Error Message	<p>If you do not use an error page, you can enter an error message to be returned to clients in the event no server is available.</p> <p>Note: This setting is not supported for SIP profiles.</p>
Security	This applies to HTTP and HTTPS only.

Settings	Guidelines
WAF Profile	Select a predefined or user-defined WAF profile configuration object. Note: This setting applies to Layer-7 HTTP/HTTPS profiles only. See Configuring a WAF Profile .
SSL Traffic Mirror	This field applies to HTTPS and TCPS only.
SSL Traffic Mirror	Select the box to enable SSL traffic mirroring and then select the ports of interest.
Traffic Log	
Log	Enable to record traffic logs for this virtual server. Note: Local logging is constrained by available disk space. We recommend that if you enable traffic logs, you monitor your disk space closely. We also recommend that you use local logging during evaluation and verification of your initial deployment, and then configure remote logging to send logs to a log management repository.
Comments	A string to describe the purpose of the configuration, to help you and other administrators more easily identify its use.

TCP multiplexing

The TCP multiplexing option enables Layer 7 load balancing virtual servers to “reuse” existing TCP connections between FortiADC and backend real servers. Using this connection pool can reduce TCP overhead and improve web server and application performance.

Figure 38: Client requests handled using connections from the connection pool



Note: The feature is not supported for profiles with the Source Address option enabled.

You can enable and configure this option using the CLI only.

To configure a connection pool and assign it to a virtual server:

Use the following command to configure the connection pool:

```
config load-balance connection-pool
  edit <name>
    set age <integer>
    set reuse <integer>
    set size <integer>
    set timeout <integer>
  next
end
```

age	Maximum duration of a connection in seconds. The recommended value is 3000.
reuse	Maximum number of times that the virtual server can reuse the connection. The recommended value is 2000.
size	Maximum number of connections in the connection pool. The recommended value is 0, which specifies that there is no limit on the connection size.
timeout	Maximum number of seconds a connection can be idle before the system deletes it. The recommended value is 30.

To assign the connection pool configuration to a virtual server, enter the following command:

```
config load-balance virtual-server
  edit <virtual-server_name>
    set type l7-load-balance
    set connection-pool <pool_name>
  end
```

where:

<pool_name> is the name of the connection pool.

Using scripts

You can use scripts to perform actions that are not currently supported by the built-in feature set. Scripts enable you to use predefined script commands and variables to manipulate the HTTP request/response or select a content route.

[Table 32](#) describes predefined scripts that you can copy and customize.

Table 32: Predefined Scripts

Predefined Script	Usage
HTTP_2_HTTPS_REDIRECTION	Redirect requests to the HTTPS site. You can use this script without changes.
REDIRECTION_by_USER_AGENT	Redirect requests based on User Agent (for example, a redirect to the mobile version of a site). You should not use this script as is. Instead, copy it and customize the User Agent and URL values.
REWRITE_HOST_n_PATH	Rewrite the host and path in the HTTP request, for example, if the site is reorganized. You should not use this script as is. Instead, copy it and customize the "old" and "new" hostnames and paths.
CONTENT_ROUTING_by_URI	Routing to a pool member based on URI string matches. You should not use this script as is. Instead, copy it and customize the URI string matches and pool member names.
CONTENT_ROUTING_by_X_FORWARDED_FOR	Routing to a pool member based on IP address in the X-Forwarded-For header. You should not use this script as is. Instead, copy it and customize the X-Forwarded-For header values and pool member names.
REDIRECTION_by_STATUS_CODE	Redirect requests based on the status code of server HTTP response (for example, a redirect to the mobile version of a site). Do NOT use this script "as is". Instead, copy it and customize the condition on the server HTTP response status code and the URL values.

You can type or paste the script content into the configuration page. After you have created a script configuration object, you can specify it in the virtual server configuration.

Before you begin:

- Create a script. See [Appendix C: Scripts](#).
- You must have Read-Write permission for System settings.

The following paragraphs shows how to:

- "Using scripts" on page 167
- "Using scripts" on page 167
- "Using scripts" on page 167
- "Using scripts" on page 167

Create a script object

To create a script configuration object:

1. Go to Server Load Balance > Scripting.
2. Click **Add** to display the configuration editor.
3. Complete the configuration as described in [Table 33](#).
4. Save the configuration.

Table 33: Script configuration

Settings	Guidelines
Name	Unique group name. Valid characters are A-Z, a-z, 0-9, __, and -. No spaces. After you initially save the configuration, you cannot edit the name.
Input	Type or paste the script.

Import a script

To import a script:

1. Click **Import**
2. Click **Choose File** to browse for the script file.
3. Click **Save**.

Export a script

To export a script:

1. Select the script of interest.
2. Click **Export**.

Delete a script

To delete a script:

1. Select the script of interest.
2. Click **Delete**.

Chapter 5: Link Load Balancing

This chapter includes the following topics:

- [Link load balancing basics](#)
- [Link load balancing configuration overview](#)
- [Configuring gateway links](#)
- [Configuring persistence rules](#)
- [Configuring proximity route settings](#)
- [Configuring a link group](#)
- [Configuring a virtual tunnel group](#)
- [Configuring link policies](#)

Link load balancing basics

The link load balancing (LLB) features are designed to manage traffic over multiple internet service provider (ISP) or wide area network (WAN) links. This enables you to subscribe to or provision multiple links, resulting in reduced risk of outages, additional bandwidth for peak events, and potential cost savings if your ISP uses billing tiers based on bandwidth rate or peak/off-peak hours.

In most cases, you configure link load balancing for outgoing traffic. Outbound traffic might be user or server traffic that is routed from your local network through your ISP transit links, leased lines, or other WAN links to destinations on the Internet or WAN. You configure link policies that select the gateway for outbound traffic.


When the FortiADC system receives outbound traffic that matches a source/destination/service tuple that you configure, it forwards it to an outbound gateway link according to system logic and policy rules that you specify.

The LLB feature supports load balancing among link groups or among virtual tunnel groups.

Using link groups

The link group option is useful for ISP links. It enables you to configure multiple ISP links that are possible routes for the traffic. The LLB picks the best route based on health checks, LLB algorithms, bandwidth rate thresholds, and other factors you specify, including a schedule.

[Figure 39](#) shows an example topology when FortiADC is deployed to support link groups.

Figure 39: LLB link groups

Using virtual tunnels

A virtual tunnel is a good choice when you want to load balance traffic from applications that embed the source address in the packet payload, like VPN and VoIP traffic. Such traffic can be difficult to load balance using traditional LLB methods. Virtual tunnels enable reliable, site-to-site connectivity using Generic Routing Encapsulation (GRE). The local FortiADC appliance encapsulates traffic so that it can be routed according to your link policy rules. The link policy rules use LLB techniques to identify the best available route among a group of links. If one of the links breaks down, the traffic can be rerouted through another link in the tunnel group. When traffic egresses the remote FortiADC appliance, it is decapsulated and the original source and destination IP addresses are restored.

Figure 40 shows an example of a deployment that does not use LLB. It uses dedicated leased lines for its WAN links, which are reliable, but expensive.

Figure 40: WAN connectivity over single leased lines





Figure 41 shows the same network deployed with FortiADC appliances. The LLB link policy load balances traffic among more affordable ADSL links.

Figure 41: LLB virtual tunnels

Depending on your business, you might use the link group option, the virtual tunnel option, or both.



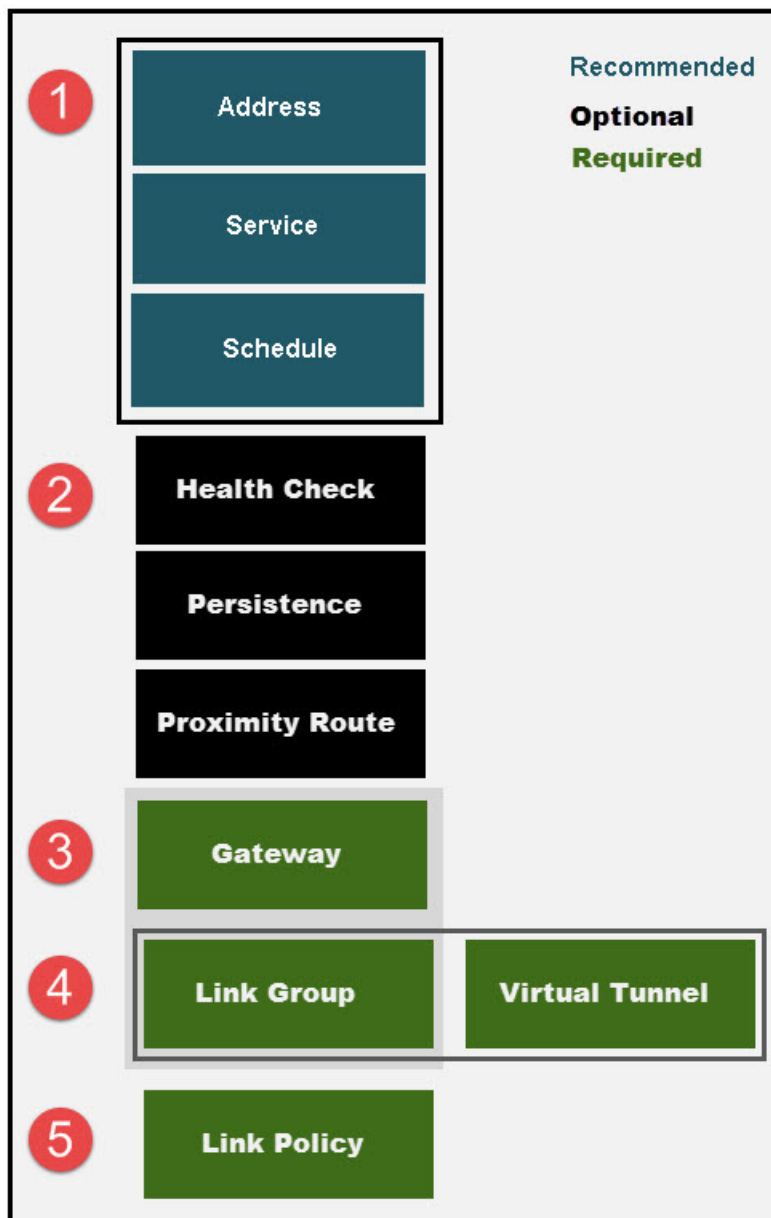
The FortiADC system evaluates traffic to determine the routing rules to apply. With regard to link load balancing, the system evaluates rules in the following order and applies the first match:

1. LLB link policy
2. Policy route
3. Static/Dynamic route
4. LLB default link group

Link load balancing configuration overview

The system has a configuration framework that enables granular link load balancing rules.

Figure 42 shows the configuration objects used in the LLB configuration and the order in which you create them. A *link policy* specifies the source/destination/service matches to which the policy applies. You apply a link policy to a *link group* or a *virtual tunnel*.

Figure 42: LLB configuration summary

The granular configuration of the gateway configuration includes health checks and bandwidth thresholds. The granular configuration of link groups includes load balancing methods, persistence rules, and proximity routes.

The granular configuration of virtual tunnels includes load balancing methods. In the virtual tunnel configuration, you can enable health check tests, but you do not use health check configuration objects.

Basic steps

1. Add address, address group, service, service group, and schedule group configuration objects that can be used to match traffic to link policy rules. This step is recommended. If your policy does not use match criteria, it will not have granularity.

2. Configure optional features. If you want to use health check rules, configure them before you configure the gateway links. If you want to use persistence rules or proximity routes, configure them before you configure a link group.
3. Configure gateway links.
4. Configure link groups or virtual tunnels.
5. Configure the link policy. When you configure a link policy, you set the source/destination/service matching tuple for your link groups or virtual tunnels.

Configuring gateway links

The gateway link configuration enables you to specify health checks, bandwidth rate thresholds, and spillover threshold behavior for the gateway links you add to link groups.

Before you begin:

- You must know the IP addresses of the ISP gateway links used in the network segment where the FortiADC appliance is deployed.
- You must have added health check configuration objects that you want to use to check the gateway links.
- You must have Read-Write permission for Link Load Balance settings.

After you have configured a gateway link configuration object, you can select it in the link group configuration.

To configure a gateway link:

1. Go to Link Load Balance > Link Group.
2. Click the **Gateway** tab.
3. Click **Add** to display the configuration editor.
4. Complete the configuration as described in [Table 34](#).
5. Save the configuration.

Table 34: LLB gateway configuration

Settings	Guidelines
Name	Configuration name. Valid characters are A-Z, a-z, 0-9, _, and -. No spaces. You reference this name in the link group configuration. Note: After you initially save the configuration, you cannot edit the name.
Address	IP address of the gateway link.
Health Check	Enable health checks.
Health Check Relationship	<ul style="list-style-type: none"> • AND—All of the selected health checks must pass for the link to be considered available. • OR—One of the selected health checks must pass for the link to be considered available.

Settings	Guidelines
Health Check List	Select one or more health check configuration objects.
Inbound Bandwidth	Maximum bandwidth rate for inbound traffic through this gateway link.
Outbound Bandwidth	<p>Maximum bandwidth rate for outbound traffic to this gateway link. If traffic exceeds this threshold, the FortiADC system considers the gateway to be full and does not dispatch new connections to it.</p> <p>The default is 2,000,000 Kbps. The valid range is 1 to 2,147,483,647.</p> <p>We recommend you tune bandwidth thresholds strategically, using the bandwidth rate and price structure agreement you have with your ISP to your advantage.</p>
Inbound Spillover Threshold	Maximum inbound bandwidth rate for a link in a spillover load balancing pool.
Outbound Spillover Threshold	<p>Maximum outbound bandwidth rate for a link in a spillover load balancing pool.</p> <p>If you enable spillover load balancing in the link group configuration, the system maintains a spillover list. It dispatches new connections to the link with the greatest priority until its spillover threshold is exceeded; then dispatches new connections to the link with the next greatest priority until its threshold is exceeded, and so on.</p> <p>The default is 2,000,000 Kbps. The valid range is 1 to 2,147,483,647.</p>
Total Spillover Threshold	Maximum total bandwidth rate (inbound plus outbound) for a link in a spillover load balancing pool.

Configuring persistence rules

Persistence rules identify traffic that should be ignored by load balancing rules and instead be forwarded to the same gateway each time the traffic traverses the FortiADC appliance.

You should use persistence rules with applications that use a secure connection. Such applications drop connections when the server detects a change in a client's source IP address.

Table 35 describes the types of persistence rules you can configure.

Table 35: Persistence rules used in link load balancing

Persistence	Description
Source-Destination Pair	Packets with the same source IP address and destination IP address take same outgoing gateway.

Persistence	Description
Source-Destination Address	Packets with a source IP address and destination IP address that belong to the same subnet take the same outgoing gateway.
Source Address	Packets with a source IP address that belongs to the same subnet take the same outgoing gateway.
Destination Address	Packets with a destination IP address that belongs to the same subnet take same outgoing gateway.

Before you begin:

- You must have an awareness of the types of outbound traffic from your network. Persistence rules are useful for traffic that requires an established session, such as secure connections (HTTPS and SSH, for example).
- You must have knowledge of the source and/or destination subnets to which the persistence rules should apply.
- You must have Read-Write permission for Link Load Balance settings.



You can use persistence rules in link groups but not virtual tunnels.

To configure a persistence rule:

1. Go to Link Load Balance > Link Group.
2. Click the **Persistence** tab.
3. Click **Add** to display the configuration editor.
4. Complete the configuration as described in [Table 36](#).
5. Save the configuration.

Table 36: Persistence rule configuration

Type	Guidelines
Name	Configuration name. Valid characters are A-Z, a-z, 0-9, _, and -. No spaces. You reference this name in the link group configuration. Note: After you initially save the configuration, you cannot edit the name.
Type	Select one of the persistence types, as described below.
Source-Destination Pair	
Timeout	The default is 300 seconds.
Source-Destination Address	

Type	Guidelines
Timeout	The default is 300 seconds.
Source IPv4 Netmask Bits	Number of bits in a subnet mask to specify a network segment that should following the persistence rule.
Destination IPv4 Netmask Bits	<p>Number of bits in a subnet mask to specify a network segment that should following the persistence rule.</p> <p>For example, if you set this to 24, and the system chooses a particular gateway router for destination IP 192.168.1.100, the system will select that same gateway for traffic to all destination IPs in subnet 192.168.1.0/24.</p>
Source Address	
Timeout	The default is 300 seconds.
Source IPv4 Netmask Bits	<p>Number of bits in a subnet mask to specify a network segment that should following the persistence rule. The default is 32, but you can set it to any value between 1 and 32.</p> <p>For example, if you set this to 24, and the system chooses a particular gateway router for client IP 192.168.1.100, the system will select that same gateway for subsequent client requests when the subsequent client belongs to subnet 192.168.1.0/24.</p>
Destination Address	
Timeout	The default is 300 seconds.
Destination IPv4 Netmask Bits	Number of bits in a subnet mask to specify a network segment that should following the persistence rule.

Configuring proximity route settings

The proximity route feature enables you to associate link groups with efficient routes. Proximity routes can improve user experience over the WAN because traffic is routed over fast routes.

You can use either or both of these methods:

- **Static Table**—You specify the gateways to use for traffic on destination networks.
- **Dynamic Detection**—The system polls the network for efficient routes. The algorithm selects a gateway based on latency.

If you configure both, the system checks the static table first for a matching route and, if any, uses it. If there is no matching static route, the system uses dynamic detection.

Before you begin:

- You must have knowledge of IP addresses used in outbound network routes to configure a static route.
- You must have Read-Write permission for Link Load Balance settings.

To configure a proximity route:

1. Go to Link Load Balance > Link Group.
2. Click the **Proximity Route** tab.
3. Complete the configuration as described in [Table 37](#).
4. Save the configuration.

Table 37: Proximity route rule configuration

Type	Guidelines
Mode	<ul style="list-style-type: none"> • Static Table First—Consult the static table first. If no match, use dynamic detection. • Static Table Only—Use the static table; do not use dynamic detection. • Dynamic Detect Only—Use dynamic detection; do not use the static table. • Disable—Do not use the proximity route configuration.
Static Table	
Type	<ul style="list-style-type: none"> • ISP—Use an ISP address object. • Subnet—Specify an IP netmask manually. <p>Routes that are specified manually have priority over ISP address object entries.</p>
ISP Name	<p>If you use the ISP configuration type, select an ISP address book configuration object.</p> <p>If an address exists in multiple ISP address books, the route entries have priority as follows:</p> <ol style="list-style-type: none"> 1. User-defined entries. 2. Entries from an address book that has been imported. 3. Entries from the predefined address book (default for the firmware image).
IP Subnet	If you use the Subnet configuration type, specify a destination IP address and netmask.
Gateway	Select a gateway configuration object. The gateway must be able to route packets to the destination IP address that you have specified.
Dynamic Detect	
Protocol	<ul style="list-style-type: none"> • ICMP—Use ICMP to detect routes. Calculate proximity by the smaller RTT. • ICMP and TCP—Some hosts do not respond to ICMP requests. Specify this option to use both ICMP and TCP to detect routes and RTT. For TCP detection, port 7 (TCP echo) is used. A connection refused or connection reset by the destination is treated as successful detection.

Type	Guidelines
Aging Period	The default is 86,400 seconds (24 hours).
Retry Number	The default is 3.
Retry Interval	The default is 3.

Configuring a link group

Link groups include ISP gateways your company uses for outbound traffic. Grouping links reduces the risk of outages and provisions additional bandwidth to relieve potential traffic congestion. See [Using link groups](#).

The link group configuration specifies the load balancing algorithm and the gateway routers in the load balancing pool. You can enable LLB options, such as persistence rules and proximity routes.

Before you begin:

- You must have configured gateway links and persistence rules and before you can select them in the link group configuration.
- You must have Read-Write permission for Link Load Balance settings.

After you have configured a link group configuration object, you can select it in the link policy configuration.

To configure a link group:

1. Go to Link Load Balance > Link Group.
The configuration page displays the Link Group tab.
2. Click **Add** to display the configuration editor.
3. Complete the configuration and add members as described in [Table 38](#).
4. Save the configuration.

Table 38: Link group configuration

Settings	Guidelines
Name	Configuration name. Valid characters are A-Z, a-z, 0-9, _, and -. No spaces. You reference this name in the LLB policy configuration. Note: After you initially save the configuration, you cannot edit the name.
Address Type	IPv4 Note: IPv4 is selected by default, and cannot be changed.

Settings	Guidelines
Route Method	<ul style="list-style-type: none"> • Weighted Round Robin—Dispatches new connections to link members using a weighted round-robin method. • Least Connections—Dispatches new connections to the link member with the lowest number of connections. • Least New Connections per Second—Dispatches new connections to the link member that has the lowest rate of new connections per second. • Least Throughput Outbound—Dispatches new connections to the link member with the least outbound traffic. • Least Throughput Inbound—Dispatches new connections to the link member with the least inbound traffic. • Least Throughput Total—Dispatches new connections to the link member with the least total traffic (that is, inbound plus outbound). • Spillover Throughput Outbound—Dispatches new connections according to the spillover list based on outbound traffic. • Spillover Throughput Inbound—Spillover list based on inbound traffic. • Spillover Throughput Total—Spillover list based on total traffic (that is, inbound plus outbound). • Source Address Hash—Selects the gateway link based on a hash of the source IP address.
Persistence	Select a persistence configuration. Optional.
Proximity Route	<ul style="list-style-type: none"> • Enable—The system uses the proximity route logic and configuration when determining routes. • Disable—The system does not use the proximity route configuration.
Add member	
Name	<p>Configuration name. Valid characters are A-Z, a-z, 0-9, _, and -. No spaces.</p> <p>After you initially save the configuration, you cannot edit the name.</p>
Gateway	Select a gateway configuration object. See Configuring gateway links .

Settings	Guidelines
Weight	<p>Assigns relative preference among members—higher values are more preferred and are assigned connections more frequently. The default is 1. The valid range is 1 to 255.</p> <p>All load balancing methods consider weight, except spillover, which uses its own priority configuration. Servers are dispatched requests proportional to their weight, relative to the sum of all weights.</p> <p>The following example shows the effect of weight on WRR:</p> <ul style="list-style-type: none"> • Sever A, Weight 2; Server B, Weight 1: Requests are sent AABAAB. • Sever A, Weight 3; Server B, Weight 2: Requests are sent AABAB. <p>For other methods, weight functions as a tie-breaker. For example, with the Least Connection algorithm, requests are sent to the server with the least connections. If the number of connections is equal, the request is sent to the server with the greater weight. For example:</p> <ul style="list-style-type: none"> • Server A, Weight 1, 1 connection • Server B, Weight 2, 1 connection <p>The next request is sent to Server B.</p>
Spillover Priority	<p>Assigns a priority to the link when using a spillover load balancing method. Higher values have greater priority. When a spillover method is enabled, the system dispatches new connections to the link that has the greatest spillover priority until its threshold is exceeded; then it dispatches new connections to the link with the next greatest priority until its threshold is exceeded, and so on.</p> <p>If multiple links in a link group have the same spillover priority, the system dispatches new connections among those links according to round robin.</p> <p>The default is 0. The valid range is 0-9.</p>
Status	<ul style="list-style-type: none"> • Enable—The member is considered available for new traffic. • Disable—The member is considered unavailable for new traffic.
Backup	<p>Enable to designate the link as a backup member of the group. All backup members are inactive until all main members are down.</p>

Configuring a virtual tunnel group

Virtual tunnels enable reliable, site-to-site connectivity using Generic Routing Encapsulation (GRE) to tunnel traffic between pairs of FortiADC appliances. See [Using virtual tunnels](#).

The virtual tunnel group configuration sets the list of tunnel members, as well as load balancing options like algorithm and weight.

When you add members to a virtual tunnel configuration, you specify a local and remote IP address. These addresses are IP addresses assigned to a network interface on the local and remote FortiADC appliance.

Before you begin:

- You must have Read-Write permission for Link Load Balance settings.

After you have configured a virtual tunnel configuration object, you can select it in the link policy configuration.

To configure a virtual tunnel:

1. Go to Link Load Balance > Virtual Tunnel.
2. Click **Add** to display the configuration editor.
3. Complete the configuration and add members as described in [Table 39](#).
4. Save the configuration.

Table 39: Virtual tunnel configuration

Settings	Guidelines
Name	Configuration name. Valid characters are A-Z, a-z, 0-9, _, and -. No spaces. You reference this name in the LLB policy configuration. Note: After you initially save the configuration, you cannot edit the name.
Method	<ul style="list-style-type: none"> • Weighted Round Robin—Dispatches packets to VT members using a weighted round-robin method. • Source-Destination Hash—Dispatches packets by source-destination IP address tuple.
Add member	
Name	Configuration name. Valid characters are A-Z, a-z, 0-9, _, and -. No spaces. After you initially save the configuration, you cannot edit the name.
Tunnel Local Address	IP address for the network interface this system uses to form a VPN tunnel with the remote system.
Tunnel Remote Address	IP address that the remote FortiADC system uses to form a VPN tunnel with this system.
Health Check	<ul style="list-style-type: none"> • Enable—Send probes to test whether the link is available. • Disable—Do not send probes to test the health of the link.
Weight	Assigns relative preference among members—higher values are more preferred and are assigned connections more frequently.
Status	<ul style="list-style-type: none"> • Enable—The member is considered available for new traffic. • Disable—The member is considered unavailable for new traffic.
Backup	Enable to designate the tunnel as a backup member of the group. All backup members are inactive until all main members are down.

Configuring link policies

A link policy matches traffic to rules that select a link group or virtual tunnel.

The policy uses a matching tuple: source, destination, service, and schedule. The policy match is a Boolean AND—All must match for the rule to be applied.

The elements of the tuple support specification by group objects. This is a Boolean OR—If source IP address belongs to member 1 OR member 2, then source matches.

The logical combinations enable you to subscribe multiple address spaces or services to a group of links, and create load balancing rules on that group basis.

The policy table is consulted from top to bottom. The first rule to match is applied.



The FortiADC system evaluates traffic to determine the routing rules to apply. With regard to link load balancing, the system evaluates rules in the following order and applies the first match:

1. LLB link policy
2. Policy route
3. Static/Dynamic route
4. LLB default link group

Before you begin:

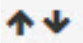
- You must have configured any address, service, and schedule objects that you want to use as match criteria for your policy.
- You must have configured a link group or virtual tunnel group.
- You must have Read-Write permission for Link Load Balance settings.

To configure a link policy:

1. Go to Link Load Balance > Link Policy.
2. Click **Add** to display the configuration editor.
3. Complete the configuration as described in [Table 40](#).
4. Save the configuration.
5. Reorder rules, as necessary.

Table 40: Link policy configuration

Option	Guidelines
Default Link Group	Select a link group configuration object that is used as the default when traffic does not match policy rules.
Name	Configuration name. Valid characters are A-Z, a-z, 0-9, _, and -. No spaces. After you initially save the configuration, you cannot edit the name.

Option	Guidelines
Ingress Interface	Select the network interface to which the policy applies.
Source Type	Whether to use address, address group, or ISP address objects for this rule.
Source, Source ISP, or Source Group	Select an address object to match source addresses. If you do not specify a source address, the rule matches any source address. See Configuring IPv4 address groups .
Destination Type	Whether to use address, address group, or ISP address objects for this rule.
Destination, Destination ISP, or Destination Group	Select an address object to match destination addresses. If you do not specify a destination address, the rule matches any destination. See Configuring IPv4 address groups .
Service Type	Whether to use service or service group objects for this rule.
Service or Service Group	Select a service object to match destination services. If you do not specify a service, the rule matches any service. See Creating service groups .
Schedule	Select the schedule object that determines the times the system uses the logic of this configuration. The link policy is active when the current time falls in a time period specified by one or more schedules in the schedule group. If you do not specify a schedule, the rule applies at all times. See Creating schedule groups .
Group Type	<ul style="list-style-type: none"> • Link Group—Policy applies to a link group. Select the option, then the link group. See Configuring a link group. • Virtual Tunnel—Policy applies to a virtual tunnel. Select the option, then the virtual tunnel. See Configuring a virtual tunnel group.
Link Group	Select a link group.
Reordering	
	After you have saved a rule, reorder rules as necessary. The rules table is consulted from top to bottom. The first rule that matches is applied and subsequent rules are not evaluated.

Chapter 6: Global Load Balancing

This chapter includes the following topics:

- [Global load balancing basics](#)
- [Global load balancing configuration overview](#)
- [Configuring data centers](#)
- [Configuring servers](#)
- [Configuring virtual server pools](#)
- [Configuring hosts](#)
- [Configuring dynamic proximity](#)
- [Configuring persistence](#)
- [Configuring an address group](#)
- [Configuring remote DNS servers](#)
- [Configuring the DSSET list](#)
- [Configuring DNS zones](#)
- [Configuring DNS64](#)
- [Configuring the response rate limit](#)
- [Configuring a Global DNS policy](#)
- [Configuring general settings](#)
- [Configuring the trust anchor key](#)

Global load balancing basics

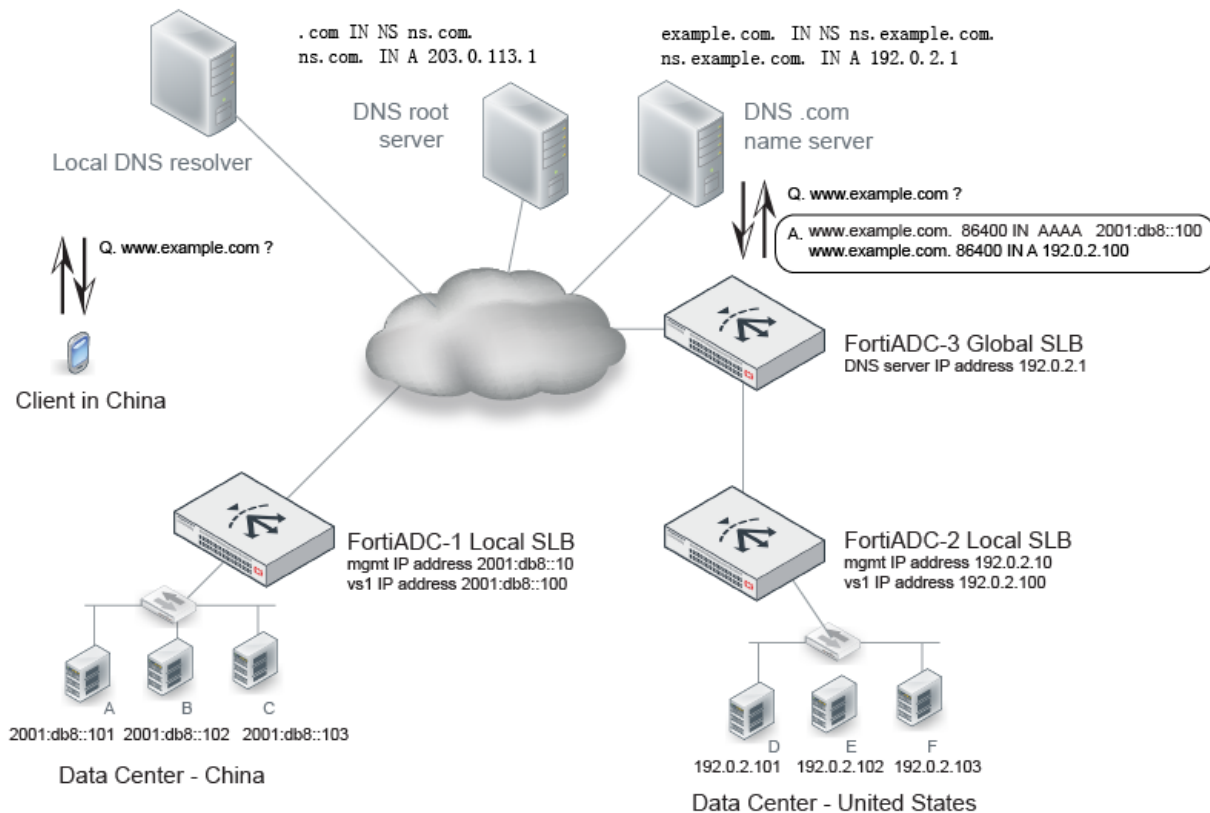
The global load balancing (GLB) feature is a DNS-based solution that enables you to deploy redundant resources around the globe that you can leverage to keep your business online when a local area deployment experiences unexpected spikes or downtime. The FortiADC system implements a hardened BIND 9 DNS server that can be deployed as the authoritative name server for the DNS zones that you configure. Zone resource records are generated dynamically based on the global load balancing framework. The DNS response to a client request is an ordered lists of answers that includes all available virtual servers. A client that receives DNS response with a list of answers tries the first and only proceeds to the next answers if the first answer is unreachable. The response list is based on the following priorities:

1. Virtual server health—Availability is determined by real-time connectivity checking. When the DNS server receives a client request, it checks connectivity for all possible matches and excludes unavailable servers from the response list.
2. Persistence—You can enable persistence for applications that have transactions across multiple hosts. A match to the persistence table has priority over proximity algorithms.
3. Geographic proximity—Proximity is determined by matching the source IP address to either the FortiGuard Geo IP database or the FortiADC predefined ISP address book.
4. Dynamic proximity—Proximity is determined by application response time (RTT probes), least connections, or byte-per-second.

5. Weighted round robin—If proximity algorithms are not configured or not applicable, available virtual servers are listed in order based on a simple load balancing algorithm.

Figure 43 shows an example global load balancing deployment with redundant resources at data centers in China and the United States.

Figure 43: Global load balancing deployment



FortiADC-1 is the local SLB for the data center in China. FortiADC-2 is the local SLB for the data center in the United States. FortiADC-3 is a global SLB. It hosts the DNS server that is authoritative for `www.example.com`. When a client clicks a link to `www.example.com`, the local host DNS resolver commences a DNS query that is ultimately resolved by the authoritative DNS server on FortiADC-3. The set of possible answers includes the virtual servers on FortiADC-1 or FortiADC-2. The global load balancing framework uses health status and proximity algorithms to determine the set of answers that are returned, and the order of the answer list. For example, you can use the global SLB framework geoproximity feature to direct clients located in China to the virtual server in China; or if the virtual server in China is unavailable, then to the redundant resources in the United States.

You configure the global load balancing framework and DNS settings only on the global FortiADC (FortiADC-3 in the example above). The virtual server IP addresses and ports can be discovered by the FortiADC global SLB from the FortiADC local SLBs. The GLB DNS server uses the discovered IP addresses in the DNS response. The framework also supports third-party IP addresses and health checks for them.

The DNS server supports the following security features:

- **DNSSEC**—Domain Name System Security Extensions. DNSSEC provides authentication by associating cryptographically generated digital signatures with DNS resource record (RR) sets. The FortiADC system makes it

easy to manage the keys that must be provided to DNS parent domains and the keys that must be imported from DNS child domains.

- Response rate limit—Helps mitigate DNS denial-of-service attacks by reducing the rate at which the authoritative name servers respond to high volumes of malicious queries.
- DNS forwarding—In a typical enterprise local area network, the client configuration has the IP address of an internal authoritative DNS server so that requests for internal resources can be answered directly from its zone data. Requests for remote resources are sent to another DNS server known as a forwarder. The internal server caches the results it learns from the forwarder, which optimizes subsequent lookups. Using forwarders reduces the number of DNS servers that must be able to communicate with Internet DNS servers.



Further reading:

BIND 9 reference manuals: <http://www.bind9.net/manuals>

RFC 1035 (DNS): <http://tools.ietf.org/html/rfc1035>

RFC 4033 (DNSSEC): <http://tools.ietf.org/html/rfc4033>

Global load balancing configuration overview

In a global load balancing deployment, you configure DNS server and global load balancing details only on the global FortiADC instance. The configuration framework enables granular administration and fine tuning of both the DNS server and the global load balancing framework.

Figure 44 shows the basic configuration elements for global load balancing and the recommended order for creating the configuration objects. The order is important for initial configurations because complex configuration elements like policies often include references to simple configuration objects like the remote DNS servers (forwarders) or DNS64 rules, but the simple elements must be created first.

Figure 44: Global load balancing configuration summary**Basic steps (DNS server)**

1. Configure address groups to use in your DNS policy matching rules. The system includes the predefined address groups **any** and **none**.
2. Configure remote DNS servers (forwarders) and the DSSET list that you might reference in the zone configuration.

3. Complete the zone configuration. The global load balancing framework generates the zone configuration for zones that include the FortiADC virtual servers.
4. Configure DNS64 or response rate limit configurations that you might reference in the DNS policy.
5. Configure the DNS policy that matches a source/destination tuple to a zone. You can also enable and configure DNSSEC in the DNS policy.
6. Configure general DNS settings to be applied when DNS requests do not match the DNS policy.

Basic steps (Global load balancing)

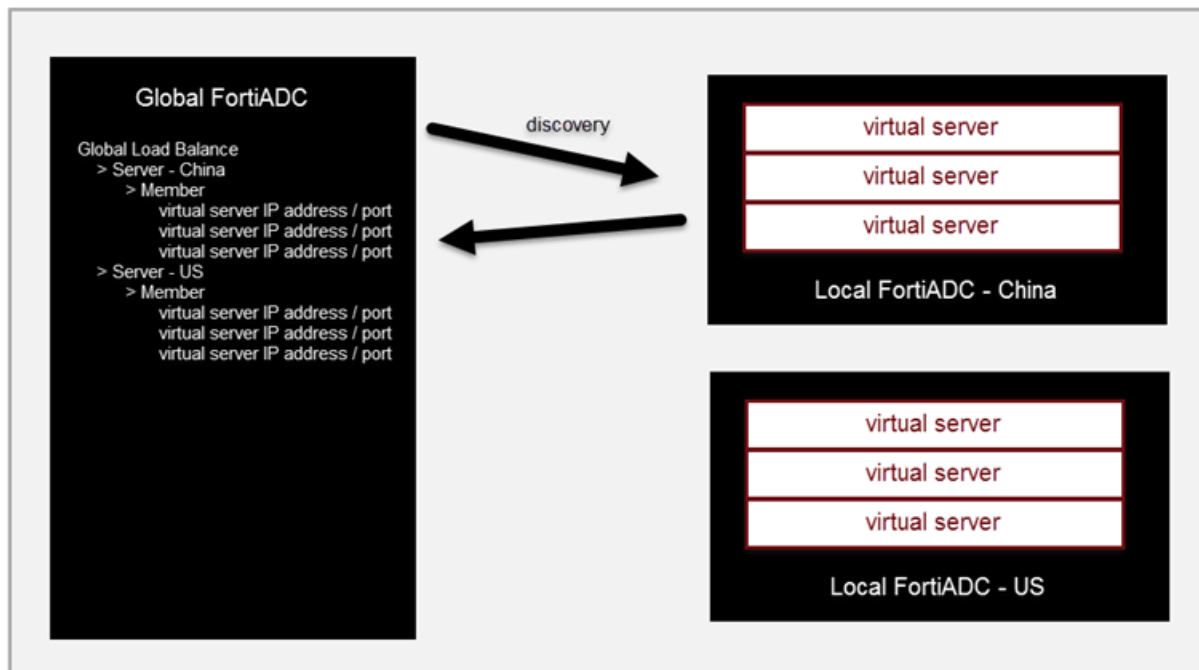
1. Create the data center, servers, virtual server pool, and host configurations that are the framework for associating locations with virtual servers and generating the DNS zone configuration and resource records. You can adjust the dynamic proximity and persistence settings at any time.
2. Review the generated DNS zone configuration.
3. Create a policy that matches traffic to the generated zone configuration.

Configuring servers

In the context of the global server load balance configuration, servers are the local SLB (FortiADC instances or third-party servers) that are to be load balanced. For FortiADC instances, the GLB checks status and synchronizes configuration from the local SLB so that it can learn the set of virtual servers that are possible to include in the GLB virtual server pool.

Figure 45 illustrates configuration discovery. Placement in this list does not include them in the pool. You also must name them explicitly in the virtual server pool configuration.

Figure 45: Virtual server discovery



Before you begin:

- You must have created the data center configuration objects that are associated with the local SLB.
- You must have created virtual server configurations on the local FortiADC SLB. In this procedure, the global SLB discovers them.
- You must have created gateway configuration objects on the local FortiADC SLB if you want to configure a gateway health check. In this procedure, the global SLB discovers them.
- You must have Read-Write permission for Global Load Balance settings.

After you have created a server configuration object, you can specify it the global load balancing virtual server pool configuration.

To configure servers:

1. Go to Global Load Balance > Global Object.
2. Click the **Server** tab.
3. Click **Add** to display the configuration editor.
4. Complete the configuration as described in [Table 41](#).
5. Use the **Discover** utility to populate the Member list configuration with virtual server configuration details from the local FortiADC SLB.
6. Optional. Edit the populated list to select a discovered gateway configuration object if you want the GSLB to perform gateway health checks.
7. Save the configuration.

Table 41: Server configuration

Settings	Guidelines
Name	Configuration name. Valid characters are A-Z, a-z, 0-9, _, and -. No spaces. You reference this name in the virtual server pool configuration. Note: After you initially save the configuration, you cannot edit the name.
Type	<ul style="list-style-type: none"> • FortiADC-SLB: A FortiADC instance. • Generic-Host: A third party ADC or server.
Synchronization	Enable/disable synchronization of the virtual server status from the local FortiADC SLB. Disabled by default. If enabled, synchronization occurs whenever there is a change in virtual server status.
Address Type	IPv4 or IPv6.
IP Address	Specify the IP address for the FortiADC management interface. This IP address is used for synchronization and also status checks. If the management interface is unreachable, the virtual servers for that FortiADC are excluded from DNS answers.
Data Center	Select a data center configuration object. The data center configuration object properties are used to establish the proximity of the servers and the client requests.

Settings	Guidelines
Health Check Control	<p>If type is Generic Host, enable/disable health checks for the virtual server list. The health check settings at this configuration level are the parent configuration. When you configure the list, you can specify whether to inherit or override the parent configuration.</p> <p>Note: This option is available only when Generic Host is selected. See Type above. Health checking is built-in, and you can optionally configure a gateway health check.</p>
Health Check Relationship	<ul style="list-style-type: none"> • AND—All of the specified health checks must pass for the server to be considered available. • OR—One of the specified health checks must pass for the server to be considered available.
Health Check List	Select one or more health check configuration objects.
Member	
Add/Delete	Add or delete member virtual servers.
Discover	Populate the member list with virtual servers from the local FortiADC configuration. After the list has been populated, you can edit the configuration to add a gateway health check.
Override	<p>Select this option if you want to update the discovered virtual server configuration with the latest configuration information whenever you use the Discover utility (for example, additions or changes to previously discovered configurations).</p> <p>Unselect this option if you want to preserve the previously discovered configuration and not have it overwritten by the Discover operation.</p>
Name	Must match the virtual server configuration name on the local FortiADC.
Address Type	IPv4 or IPv6.
IP Address	Virtual server IP address.
Port	Virtual server port.
Protocol	TCP or UDP. The default is TCP.

Settings	Guidelines
Gateway	<p>Enable an additional health check: is the gateway beyond the FortiADC reachable?</p> <p>The list of gateway configuration objects is populated by discovery, but you must select the appropriate one from the list.</p>
Health Check Inherit	If type is Generic Host, enable to inherit the health check settings from the parent configuration. Disable to specify health check settings in this member configuration.
Health Check Control	<p>Enable health checking for the virtual server.</p> <p>Note: This option is available only when Health Check Inherit is disabled. In that case, you can enable this option and configure the Health Check Relationship and Health Check List fields below.</p>
Health Check Relationship	<ul style="list-style-type: none"> • AND—All of the specified health checks must pass for the server to be considered available. • OR—One of the specified health checks must pass for the server to be considered available.
Health Check List	Specify one or more health check configuration objects.

Configuring a global load balance link

To configure a global load balance link:

1. Go to Global Load Balance > Global Object.
2. Click the **Link** tab.
3. Click **Add** to display the configuration editor.
4. Complete the configuration as described in [Table 42](#).
5. Save the configuration.

Table 42: Global load balance link configuration

Settings	Guidelines
Link	
Name	<p>Configuration name. Valid characters are A-Z, a-z, 0-9, _, and -. No spaces. You reference this name in the global load balance servers configuration.</p> <p>Note: After you initially save the configuration, you cannot edit the name.</p>

Settings	Guidelines
Data Center	Select a data center from the list. Note: You must the data center(s) configured ahead of time.
ISP	Select an ISP from the list.
ISP Province	Select an ISP province from the list.
Gateway	
Server	Select a server.
Gateway Name	Specify a name for the gateway.
or Select Here	Click the down arrow to select a gateway from the drop-down list. Note: Use this option only when you already have a list of gateways configured on the system.

Configuring data centers

The data center configuration sets key properties: Location and/or ISP and ISP province. These properties are used in the global load balancing algorithm that selects the FortiADC in closest proximity to the client.

Before you begin:

- If you want to select a user-defined ISP address book, you must create it before creating the data center configuration.
- You must have Read-Write permission for Global Load Balance settings.

After you have created a data center configuration object, you can specify it in the global load balance servers configuration.

To configure a data center:

1. Go to Global Load Balance > Global Object.
2. Click the **Data Center** tab.
3. Click **Add** to display the configuration editor.
4. Complete the configuration as described in [Table 43](#).
5. Save the configuration.

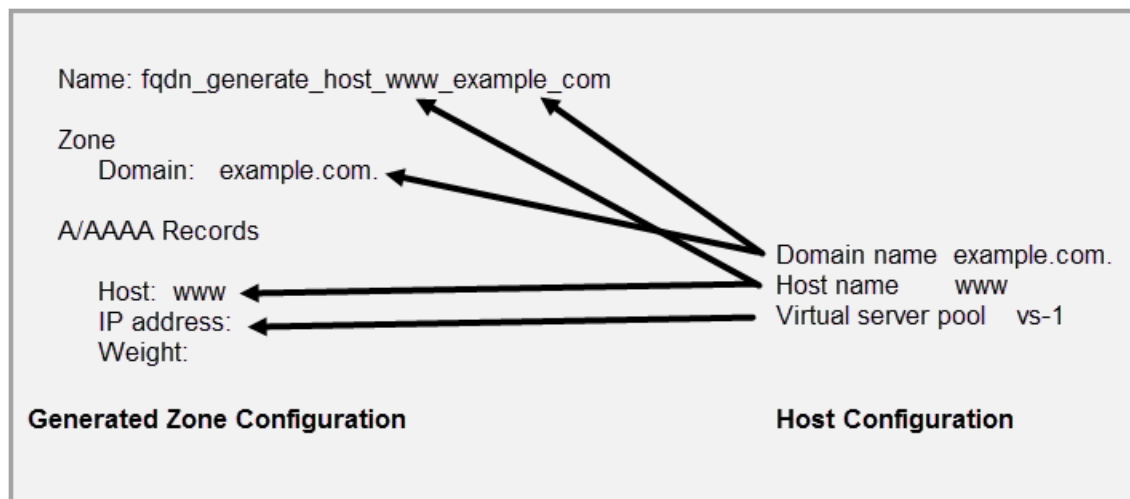
Table 43: Data center configuration

Settings	Guidelines
Name	Configuration name. Valid characters are A-Z, a-z, 0-9, _, and -. No spaces. You reference this name in the global load balance servers configuration. Note: After you initially save the configuration, you cannot edit the name.
Location	Select a location from the location list.
Description	Optional description to help administrators know the purpose or usage of the configuration.

Configuring hosts

Host settings are used to form the zone configuration and resource records in the generated DNS zone used for global load balancing.

Figure 46 shows how the host settings are mapped to zone settings and resource records. Domain and hostname are used in both the configuration and the generated configuration name. The IP address and weight are derived from the virtual server pool.

Figure 46: Host configuration and the generated DNS zone

Before you begin:

- You must have created the global virtual server pools you want to use.
- You must have Read-Write permission for Global Load Balance settings.

After you have created a host configuration object, it can be used to form the zone and resource records in the generated DNS zone configuration.

To configure a host:

1. Go to Global Load Balance > FQDN Settings.
2. Click the **Host** tab.
3. Click **Add** to display the configuration editor.
4. Complete the configuration as described in [Table 44](#).
5. Save the configuration.

Table 44: Host configuration

Settings	Guidelines
Name	<p>Configuration name. Valid characters are A-Z, a-z, 0-9, _, and -. No spaces.</p> <p>Note: After you initially save the configuration, you cannot edit the name.</p>
Host Name	<p>The hostname part of the FQDN, such as <code>www</code>.</p> <p>Note: You can specify the @ symbol to denote the zone root. The value substituted for @ is the preceding \$ORIGIN directive.</p>
Domain Name	<p>The domain name must end with a period. For example: <code>example.com.</code></p>
Respond Single Record	<p>Enable/disable an option to send a single record in response to a query. Disabled by default. By default, the response is an ordered list of records.</p>
Persistence	<p>Enable/disable the persistence table. Disabled by default.</p> <p>If you enable persistence, the client source address is recorded in the persistence table, and subsequent requests from the same network or the same host or domain are sent an answer with the virtual servers listed in the same order (unless a server becomes unavailable and is therefore omitted from the answer).</p>
Load Balance Method	<ul style="list-style-type: none"> • None • Topology—If selected, virtual servers with the same topology information as the local DNS address will be responded. • Global Availability—If selected, virtual servers will be responded by their global availability: the first virtual server in queue will always be responded if it is globally available, and the next virtual server in queue will be responded if the preceding virtual server is unavailable.
Default Feedback IPv4	<p>Specify an IP address to return in the DNS answer if no virtual servers are available.</p>
Default Feedback IPv6	<p>Specify an IPv6 address to return in the DNS answer if no virtual servers are available.</p>
Virtual Server Pool	

Settings	Guidelines
Name	Enter the mkey.
Virtual Server Pool	Select a virtual server pool from the list, or create a new one.
Weight	Assign a weight. Valid values range from 1 to 255.
Topology	Select a topology from the list, or create a new one.
ISP	Select an ISP from the list or create a new one.

Configuring virtual server pools

The virtual server pool configuration defines the set of virtual servers that can be matched in DNS resource records, so it should include, for example, all the virtual servers that can be answers for DNS requests to resolve `www.example.com`.

You also specify the key parameters of the global load balancing algorithm, including proximity options, status checking options, load balancing method, and weight.

The DNS response is an ordered list of answers. Virtual servers that are unavailable are excluded. Available virtual servers are ordered based on the following priorities:

1. Geographic proximity
2. Dynamic proximity
3. Weighted round robin

A client that receives DNS response with a list of answers tries the first and only proceeds to the next answers if the first answer is unreachable.

Before you begin:

- You must have created GLB Servers configuration objects.
- You must have Read-Write permission for Global Load Balance settings.

After you have created a virtual server pool configuration object, you can specify it in the global load balancing host configuration.

To configure a virtual server pool:

1. Go to Global Load Balance > FQDN Settings.
2. Click the **Virtual Server Pool** tab.
3. Click **Add** to display the configuration editor.
4. Complete the configuration as described in [Table 45](#).
5. Save the configuration.

Table 45: Virtual server pool configuration

Settings	Guidelines
Name	<p>Specify a unique name for the virtual server pool configuration. Valid characters are A-Z, a-z, 0-9, _, and -. No spaces. You reference this name in the host configuration.</p> <p>Note: After you initially save the configuration, you cannot edit the name.</p>
Preferred	<ul style="list-style-type: none"> • None—No preference. • Geo—If selected, virtual servers with the same GEO information as the local DNS address will be responded. • Geo-ISP—If selected, virtual servers with the same ISP information as the local DNS address will be responded first, and virtual servers with the same GEO information as the local DNS address will be responded second. • RTT—Virtual servers with the shortest latency link or closest to the data center will be responded. • Least-Connections—Virtual servers with the least connections will be responded. • Connection-Limit—Virtual servers will be responded by their connection limit determined by virtual servers' weight: the greater the weight of a virtual server, the more responses it will get. • Bytes-Per-Second—Virtual servers with the lowest traffic will be responded.
Alternate	Same as above.
Load Balance Method	Weighted Round Robin
Check Server Status	Enable/disable polling of the local FortiADC SLB. If the server is unresponsive, its virtual servers are not selected for DNS answers.
Check Virtual Server Existence	Enable/disable checks on whether the status of the virtual servers in the virtual server list is known. Virtual servers with unknown status are not selected for DNS answers.
Member	
Server	Select a GLB Servers configuration object.
Server Member	Select the name of the virtual server that is in the servers virtual server list configuration.

Settings	Guidelines
Weight	Assigns relative preference among members—higher values are more preferred and are assigned connections more frequently. The default is 1. The valid range is 1-255.
Backup	Enable to designate the member as a backup. Backup members are inactive until all main members are down.

Configuring Topologies

Before you begin:

- You must have Read-Write permission for Global Load Balance settings.

To configure a topology:

- Go to Global Load Balance > FQDN Settings.
- Click the **Topology** tab.
- Click **Add** to display the configuration editor.
- Complete the configuration as described in [Table 46](#).
- Save the configuration.

Table 46: Topology configuration guideline

Settings	Guidelines
Name	Specify the name of the topology configuration. Valid characters are A-Z, a-z, 0-9, _, and -. No spaces. After you initially save the configuration, you cannot edit the name.
GEO IP List	Select the geo IPs from the Available Items list and add them to the Selected Items list.

Configuring dynamic proximity

Use this page to configure dynamic proximity. Dynamic proximity is used to order DNS lookup results based on round-trip time (RTT) for ICMP or TCP probes sent by the local SLB to the DNS resolver that sent the DNS request.

The system caches the RTT results for the period specified by the timeout. When there are subsequent requests from clients that have a source IP address within the specified netmask, the RTT is taken from the results table instead of a new, real-time probe. This reduces DNS response time.

Before you begin:

- You must have Read-Write permission for Global Load Balance settings.

The settings you configure are applied if dynamic-proximity is enabled in the virtual server pool configuration.

To configure dynamic proximity settings:

1. Go to Global Load Balance > FQDN Settings.
2. Click the **GLB Proximity** tab.
3. Complete the configuration as described in [Table 47](#).
4. Save the configuration.

Table 47: Dynamic proximity settings

Settings	Guidelines
Protocol	<ul style="list-style-type: none"> • ICMP—Use ICMP to detect routes. Calculate proximity by the smaller RTT. • ICMP and TCP—Some hosts do not respond to ICMP requests. Specify this option to use both ICMP and TCP to detect routes and RTT. For TCP detection, a SYN packet is sent to port 53. A connection refused or connection reset by the destination is treated as successful detection.
Retry Number	Retry count if the probe fails. The default is 3. The valid range is 1-10 times.
Retry Interval	Interval between retries if the probe fails. The default is 3. The valid range is 1-3600 seconds.
IPv4 Prefix Length	Number of IPv4 netmask bits that define network affinity for the RTT table. The default is 24. For example, if the GLB records an RTT for a client with source IP address 192.168.1.100, the record is stored and applies to all requests from the 192.168.1.0/24 network.
IPv6 Prefix Length	Number of IPv6 netmask bits that define network affinity for the RTT table. The default is 64.
Aging Timeout	RTT results are cached. This setting specifies the length of time in seconds for which the RTT cache entry is valid. The default is 86400. The valid range is 60-2,592,000 seconds.

Configuring persistence

Use this page to configure source address affinity and a timeout for GSLB persistence. You enable persistence per host in the GSLB host configuration.

If the DNS query is for a host that has persistence enabled, the DNS server replies with an answer that has the virtual server IP addresses listed in the order determined by the GSLB proximity algorithms, and the client source IP address (for example 192.168.1.100) is recorded in the persistence table. If source address affinity is set to 24 bits, subsequent queries for the host from the 192.168.1.0/24 network are sent an answer with the virtual servers listed in the same order (unless a server becomes unavailable and is therefore omitted from the answer).

Persistence is required for applications that include transactions across multiple hosts, so the persistence table is also used for queries for other hosts with the same domain. For example, a transaction on a banking application

might include connections to login.bank.com and transfer.bank.com. To support persistence in these cases, the GSLB persistence lookup accounts for domain as well. The first query for login.bank.com creates a mapping for the source address network 192.168.1.0/24 and the domain bank.com. When the DNS server receives subsequent requests, it consults the persistence table for a source network match, then a domain match and a hostname match. In this example, as long as you have created host configurations for both login.bank.com and transfer.bank.com, and persistence is enabled for each, the persistence table can be used to ensure the DNS answers to queries from the same network list the resource records in the same order.

Before you begin:

- You must have Read-Write permission for Global Load Balance settings.

To configure persistence:

1. Go to Global Load Balance > FQDN Settings.
2. Click the **Persistence** tab.
3. Complete the configuration as described in [Table 48](#).
4. Save the configuration.

Table 48: GSLB persistence settings

Settings	Guidelines
IPv4 Mask Length	Number of IPv4 netmask bits that define network affinity for the persistence table. The default is 24.
IPv6 Mask Length	Number of IPv6 netmask bits that define network affinity for the persistence table. The default is 64.
Aging Period	This setting specifies the length of time in seconds for which the entry is maintained in the persistence table. The default is 86400. The valid range is 60-2,592,000 seconds.

Configuring an address group

An address group is a configuration object that specifies the source and destination IP addresses that are the matching criteria for DNS policies.

Before you begin:

- You must have Read-Write permission for Global Load Balance settings.

After you have configured an address group, you can select it in the DNS policy configuration.

To configure address groups:

1. Go to Global Load Balance > Zone Tools.
2. Click the **Address Group** tab.
3. Click **Add** to display the configuration editor.
4. Complete the configuration and add members as described in [Table 49](#)

Table 49: Address group configuration

Settings	Guidelines
Name	Configuration name. Valid characters are A-Z, a-z, 0-9, _, and -. No spaces. You reference this name in the global DNS policy configuration. Note: After you initially save the configuration, you cannot edit the name.
Member	
Address Type	<ul style="list-style-type: none"> • IPv4 • IPv6
IP/Netmask	Address/mask notation to match the IP address in the packet header. Create objects to match source IP address and different objects to match destination IP address.
Action	<ul style="list-style-type: none"> • Include—The rule logic creates an address object that includes addresses matching the specified address block. • Exclude—The rule logic creates an address object that excludes addresses matching the specified address block.

Configuring remote DNS servers

The remote server configuration is used to create a list of DNS forwarders. DNS forwarders are commonly used when you do not want the local DNS server to connect to Internet DNS servers. For example, if the local DNS server is behind a firewall and you do not want to allow DNS through that firewall, you implement DNS forwarding to a remote server that is deployed in a DMZ or similar network region that can contact Internet DNS servers.

Before you begin:

- You must have a good understanding of DNS and knowledge of the remote DNS servers that can be used to communicate with Internet domain servers.
- You must have Read-Write permission for Global Load Balance settings.

After you have configured remote DNS servers, you can select them in DNS zone and DNS policy configurations.

To configure a remote server:

1. Go to Global Load Balance > Zone Tools.
2. Click the **Remote DNS Server** tab.
3. Click **Add** to display the configuration editor.
4. Complete the configuration and add members as described in [Table 50](#).

Table 50: Remote DNS server configuration

Settings	Guidelines
Name	Configuration name. Valid characters are A-Z, a-z, 0-9, _, and -. No spaces. You reference this name in the zone configuration (if you use forwarders). Note: After you initially save the configuration, you cannot edit the name.
Member	
Address Type	<ul style="list-style-type: none"> IPv4 IPv6
Address	IP address of the remote DNS server.
Port	Port number the remote server uses for DNS. The default is 53.

Configuring the DSSET list

If you enable DNSSEC, secure communication between the FortiADC DNS server and any child DNS servers is based on keys contained in delegation signer files (DSSET files). In DNSSEC deployments, DSSET files are generated automatically when the zone is signed by DNSSEC.

You use the DSSET list configuration to paste in the content of the DSSET files provided by child domain servers or stub domains.

Note: You use the Global DNS zone configuration to generate the DSSET file for this server. The file generated by the zone configuration editor is the one you give to any parent zone or the registrar of your domain.

Before you begin:

- You must have a good understanding of DNSSEC and knowledge of the DNS deployment in your network.
- You must have used DNSSEC to sign the child domain servers and have downloaded the DSset files to a location you can reach from your management computer.
- You must have Read-Write permission for Global Load Balance settings.

After you have configured a DSSET list, you can select it in DNS zone configuration.

To configure the DSSET list:

- Go to Global Load Balance > Zone Tools.
- Click the **DSSET List** tab.
- Click **Add** to display the configuration editor.
- Complete the configuration as described in [Table 51](#).

Table 51: DSset list configuration

Settings	Guidelines
Name	Configuration name. Valid characters are A-Z, a-z, 0-9, _, and -. No spaces. You reference the name in the zone configuration (if you enable DNSSEC). After you initially save the configuration, you cannot edit the name.
Filename	Type the filename. The convention is dsset-<domain>, for example, dsset-example.com.
Content	Paste the DSset file content. The content of DSset files is similar to the following: <pre>dns.example.com. IN DS 13447 5 1 A5AD9EFB6840F58CF817F3CC7C24A7ED2DD5559C</pre>

Configuring DNS zones

The DNS zone configuration is the key to the global load balancing solution. This configuration contains the key DNS server settings, including:

- Domain name and name server details.
- Type—Whether the server is the master or a forwarder.
- DNSSEC—Whether to use DNSSEC.
- DNS RR records—The zone configuration contains resource records (RR) used to resolve DNS queries delegated to the domain by the parent zone.

You can specify different DNS server settings for each zone you create. For example, the DNS server can be a master for one zone and a forwarder for another zone.

Before you begin:

- You must have a good understanding of DNS and knowledge of the DNS deployment in your network.
- You must have authority to create authoritative DNS zone records for your network.
- You must have Read-Write permission for Global Load Balance settings.

After you have configured a DNS zone, you can select it in the DNS policy configuration.

To configure the DNS zone:

1. Go to Global Load Balance > Zone Tools.
2. Click the **Zone** tab.
3. Click **Add** to display the configuration editor.
4. Complete the configuration as described in [Table 52](#).

Table 52: DNS zone configuration

Settings	Guidelines
Name	<p>Configuration name. Valid characters are A-Z, a-z, 0-9, <code>_</code>, and <code>-</code>. No spaces. You reference the name in the global DNS policy configuration.</p> <p>Note:</p> <ul style="list-style-type: none"> FortiADC supports third-party domain names. After you initially save the configuration, you cannot edit the name.
Type	<ul style="list-style-type: none"> Master—The configuration contains the “master” copy of data for the zone and is the authoritative server for it. Forward—The configuration allows you to apply DNS forwarding on a per-domain basis, overriding the forwarding settings in the “general” configuration. FQDN Generate—The zone and its resource record is generated from the global load balancing framework.
Domain Name	The domain name must end with a period. For example: <code>example.com.</code>
Forward Options	
Forward	<ul style="list-style-type: none"> First—The DNS server queries the forwarder before doing its own DNS lookup. Only—Only query the forwarder. Do not perform a DNS lookup. Note: The internal server caches the results it learns from the forwarders, which optimizes subsequent lookups.
Forwarders	Select a remote server configuration object.
Master Options	
TTL	<p>The <code>\$TTL</code> directive at the top of the zone file (before the SOA) gives a default TTL for every RR without a specific TTL set.</p> <p>The default is 86,400. The valid range is 0 to 2,147,483,647.</p>
Negative TTL	The last field in the SOA—the negative caching TTL. This informs other servers how long to cache no-such-domain (NXDOMAIN) responses from you. The default is 3600 seconds. The valid range is 0 to 2,147,483,647.
Responsible Mail	<p>Username of the person responsible for this zone, such as <code>hostmaster.example.com.</code></p> <p>Note: Format is <code>mailbox-name.domain.com.</code> (remember the trailing dot). The format uses a dot, not the <code>@</code> sign used in email addresses because <code>@</code> has other uses in the zone file. Email, however, is sent to <code>hostmaster@example.com</code>.</p>
Primary Server Name	Sets the server name in the SOA record.

Settings	Guidelines
Primary Server Address	The IP address of the primary server.
DNSSEC	Enable/disable DNSSEC.
DNSSEC Algorithm	Only RSASHA1 is supported.
KSK Filename	It is generated by the system if DNSSEC is enabled for the zone. To regenerate the KSK, disable DNSSEC and then re-enable DNSSEC.
KSK	Type characters for a string key. The file is generated by the system if DNSSEC is enabled for the zone.
ZSK Filename	It is generated by the system if DNSSEC is enabled for the zone. To regenerate the ZSK, disable DNSSEC and then re-enable DNSSEC.
ZSK	Type characters for a string key. The file is generated by the system if DNSSEC is enabled for the zone.
DSSET Filename	The file is generated by the system if DNSSEC is enabled for the zone. The file generated by the zone configuration editor is the one you give to any parent zone or the registrar of your domain. The convention is dsset-<domain>, for example <code>dsset-example.com</code> .
DSSET	It is generated by the system if DNSSEC is enabled for the zone.
DSSET List	Select a DSSET configuration object. See Configuring the DSSET list .
FQDN Record	
FQDN Record table	Displays a summary of all DNS RR for the zone, including generated and manually configured RR.
A/AAAA Record	
Hostname	The hostname part of the FQDN, such as <code>www</code> . Note: You can specify the @ symbol to denote the zone root. The value substituted for @ is the preceding \$ORIGIN directive.
Type	<ul style="list-style-type: none"> • IPv4 • IPv6

Settings	Guidelines
Weight	Assigns relative preference among members—higher values are more preferred and are assigned connections more frequently. The default is 1. The valid range is 1-255.
Address	Specify the IP address of the virtual server.
Method	Weighted Round Robin is the only method supported.
CNAME Record	
Alias	An alias name to another true or canonical domain name (the target). For instance, <code>www.example.com</code> is an alias for <code>example.com</code> .
Target	The true or canonical domain name. For instance, <code>example.com</code> .
NS Record	
Domain Name	The domain for which the name server has authoritative answers, such as <code>example.com</code> . Note: FortiADC supports third-party domain names.
Hostname	The hostname part of the FQDN, such as <code>ns</code> .
Type	<ul style="list-style-type: none"> • IPv4 • IPv6
Address	Specify the IP address of the name server.
MX Record	
Hostname	The hostname part of the FQDN for a mail exchange server, such as <code>mail</code> .
Priority	Preference given to this RR among others at the same owner. Lower values have greater priority.
Type	<ul style="list-style-type: none"> • IPv4 • IPv6
Address	Specify the IP address.
TXT Record	

Settings	Guidelines
Name	<p>Hostname.</p> <p>TXT records are name-value pairs that contain human readable information about a host. The most common use for TXT records is to store SPF records.</p>
Text	<p>Comma-separated list of name=value pairs.</p> <p>An example SPF record has the following form:</p> <pre>v=spf1 +mx a:colo.example.com/28 -all</pre> <p>If you complete the entry from the the Web UI, do not put the string in quotes. (If you complete the entry from the CLI, you do put the string in quotes.)</p>
SRV Record	
Host Name	The host name part of the FQDN, e.g., <code>www</code> .
Priority	A priority assigned to the target host: the lower the value, the higher the priority.
Weight	A relative weight assigned to a record among records of the same priority: the greater the value, the more weight it carries.
Port	The TCP or UDP port on which the service is provided.
Target Name	The canonical name of the machine providing the service.

Configuring DNS64

The DNS64 configuration maps IPv4 addresses to AAAA queries when there are no AAAA records. This feature is optional. It can be used in network segments that use NAT64 to support IPv6 client communication with IPv4 backend servers.

Before you begin:

- You must have a good understanding of DNS and knowledge of the DNS deployment in your network.
- You must have configured address objects that specify the network segments for which the DNS64 map applies. See [Configuring an address group](#).
- You must have Read-Write permission for Global Load Balance settings.

After you have created a DNS64 configuration, you can select it a DNS policy configuration.

To configure DNS64:

1. Go to Global Load Balance > Zone Tools.
2. Click the **DNS64** tab.
3. Click **Add** to display the configuration editor.
4. Complete the configuration as described in [Table 53](#).

Table 53: DNS64 configuration

Settings	Guidelines
Name	Configuration name. Valid characters are A-Z, a-z, 0-9, _, and -. No spaces. You reference the name in the global DNS policy configuration. After you initially save the configuration, you cannot edit the name.
IPv6 Prefix	IP address and netmask that specify the DNS64 prefix. Compatible IPv6 prefixes have lengths of 32, 40, 48, 56, 64 and 96 as per RFC 6052. Each DNS64 configuration has one prefix. Multiple configurations can be defined.
Source Address	Select an address object. Only clients that match the source IP use the DNS64 lookup table.
Mapped Address	Select an address object that specifies the IPv4 addresses that are to be mapped in the corresponding A RR set.
Exclude	Select an address object. Allows specification of a list of IPv6 addresses that can be ignored. Typically, you exclude addresses that do have AAAA records.

Configuring the response rate limit

The response rate limit keeps the FortiADC authoritative DNS server from being used in amplifying reflection denial of service (DoS) attacks.

Before you begin:

- You must have a good understanding of DNS.
- You must have Read-Write permission for Global Load Balance settings.

After you have created a response rate limit configuration, you can select it in the DNS policy and DNS general settings configurations.

To configure the response rate limit:

1. Go to Global Load Balance > Zone Tools.
2. Click the **Response Rate Limit** tab.
3. Click **Add** to display the configuration editor.
4. Complete the configuration as described in [Table 54](#).

Table 54: Response rate limit configuration

Settings	Guidelines
Name	Configuration name. Valid characters are A-Z, a-z, 0-9, _, and -. No spaces. You reference the name in the global DNS policy configuration. After you initially save the configuration, you cannot edit the name.
Responses per Second	Maximum number of responses per second. The valid range is 1-2040. The default is 1000.

Configuring a Global DNS policy

The Global DNS policy is a rule base that matches traffic to DNS zones. Traffic that matches both source and destination criteria is served by the policy. Traffic that does not match any policy is served by the DNS “general settings” configuration.

Before you begin, you must have:


- A good understanding of DNS and knowledge of the DNS deployment in your network.
- Configured address objects, remote servers, DNS zones, and optional configuration objects you want to specify in your policy.
- Read-Write permission for Global Load Balance settings.

To configure the global DNS policy rule base:

1. Go to Global Load Balance > Zone Tools.
2. Click the **Global DNS Policy** tab.
3. Click **Add** to display the configuration editor.
4. Complete the configuration as described in [Table 55](#).
5. Save the configuration.
6. Reorder rules, as necessary.

Table 55: Global DNS policy configuration

Settings	Guidelines
Name	Configuration name. Valid characters are A-Z, a-z, 0-9, _, and -. No spaces. After you initially save the configuration, you cannot edit the name.
Source	Select an address object to specify the source match criteria. See Configuring an address group .
Destination	Select an address object to specify the destination match criteria. See Configuring an address group .

Settings	Guidelines
Zone List	Select one or more zone configurations to serve DNS requests from matching traffic. See Configuring DNS zones .
DNS64 List	Select one or more DNS64 configurations to use when resolving IPv6 requests. See Configuring DNS64 .
Recursion	Enables/disables recursion. If enabled, the DNS server attempts to do all the work required to answer the query. If not enabled, the server returns a referral response when it does not already know the answer.
DNSSEC	Enables/disables DNSSEC.
DNSSEC Validation	Enables/disables DNSSEC validation.
Forward	<ul style="list-style-type: none"> • First—The DNS server queries the forwarders list before doing its own DNS lookup. • Only—Only queries the forwarders list. Does not perform its own DNS lookups. <p>Note: The internal server caches the results it learns from the forwarders, which optimizes subsequent lookups.</p>
Forwarders	If the DNS server zone has been configured as a forwarder, select the remote DNS server to which it forwards requests. See Configuring remote DNS servers .
Response Rate Limit	Select a rate limit configuration object. See Configuring the response rate limit .
Reordering	
	After you have saved a rule, reorder rules as necessary. The rules table is consulted from top to bottom. The first rule that matches is applied and subsequent rules are not evaluated.

Configuring general settings

The general settings configuration specifies the interfaces that listen for DNS requests. By default, the system listens on the IPv4 and IPv6 addresses of all configured interfaces for DNS requests.

The other settings in the general settings configuration are applied when traffic does not match a Global DNS policy.

Before you begin:

- You must have a good understanding of DNS and knowledge of the DNS deployment in your network.
- You must have Read-Write permission for Global Load Balance settings.

To configure general settings:

1. Go to Global Load Balance > Zone Tools.
2. Click the **General Settings** tab.
3. Complete the configuration as described in [Table 56](#).
4. Save the configuration.

Table 56: General configuration

Settings	Guidelines
Global DNS Configuration	Enables/disables this configuration.
Recursion	Enables/disables recursion. If enabled, the DNS server attempts to do all the work required to answer the query. If not enabled, the server returns a referral response when it does not already know the answer.
DNSSEC	Enables/disables DNSSEC.
DNSSEC Validation	Enables/disables DNSSEC validation.
Listen on IPv6	Enables/disables listening for DNS requests on the interface IPv6 address.
Listen on IPv4	Enables/disables listening for DNS requests on the interface IPv4 address.
Traffic Log	Enables/disables traffic log.
Listen on All Interface	Enables listening on all interfaces.
Forward	<ul style="list-style-type: none"> • First—The DNS server queries the forwarder before doing its own DNS lookup. • Only—Only queries the forwarder. Does not perform its own DNS lookups. <p>Note: The internal server caches the results it learns from forwarders, which optimizes subsequent lookups.</p>
Use System DNS Server	Forwards DNS requests to the system DNS server instead of the forwarders list.
Response Rate Limit	Selects a rate limit configuration object. See Configuring the response rate limit .

Configuring the trust anchor key

DNSSEC validation requires that a DNS name server know the trust anchor key for the root DNS domain in order to validate already signed responses. In general, trust anchor keys do not change often, but they do change occasionally, and might change unexpectedly in the event the keys are compromised.

The FortiADC DNS server is preconfigured with a trust anchor key for the root DNS domain. If you are informed that you must update this key, you can use the configuration editor to paste the new content into the DNS server configuration.

Further reading:

<http://data.iana.org/root-anchors/draft-icann-dnssec-trust-anchor.html>

Before you begin:

- You must have a good understanding of DNSSEC and knowledge of the DNS deployment in your network.
- You must have already obtained the key so that you can copy and paste it into the DNS server configuration.
- You must have Read-Write permission for Global Load Balance settings.

To configure the trust anchor key:

1. Go to Global Load Balance > Zone Tools.
2. Click the **Trust Anchor Key** tab.
3. Click **Add** to display the configuration editor.
4. Complete the configuration as described in [Table 57](#).
5. Save the configuration.

Table 57: Trust anchor key configuration

Settings	Guidelines
Name	Configuration name. Valid characters are A-Z, a-z, 0-9, _, and -. No spaces. After you initially save the configuration, you cannot edit the name.
Value	The key value. The key format is a string with the following format: <code>\ "<domainname>" <num1> <num2> <num3> \ "<content>"</code> The following is an example: <code>\ ".\" 256 3 5 \"AwEAAbDrWmiIReotvZ6F0bgKygzWUxSUJW9z5pjiQMLH0JBGXooHrR16 pdKhI9mNkM8bLUMtwYfgeUOYXIvfagee8rk=\"</code>
Description	Description for the key.

Chapter 7: Network Security

This chapter includes the following topics:

- Security features basics
- Managing IP reputation policy settings
- Configuring IP reputation exception
- Using the Geo IP block list
- Enabling denial of service protection
- Configuring a firewall policy
- Configuring firewall connection limit

Security features basics

In most deployment scenarios, we recommend you deploy FortiGate to secure your network. Fortinet includes security functionality in the FortiADC system to support those cases when deploying FortiGate is impractical. FortiADC includes the following security features:

- Firewall—Drop traffic that matches a source/destination/service tuple you specify.
- Security connection limit—Drop an abnormally high volume of traffic from a source/destination/service match.
- IP Reputation service—Drop or redirect traffic from source IPs that are on the FortiGuard IP Reputation list.
- Geo IP—Drop or redirect traffic from source IPs that correspond with countries in the FortiGuard Geo IP database.
- Web application firewall—Drop or alert when traffic matches web application firewall attack signatures and heuristics.
- Denial of service protection—Drop half-open connections to protect the system from a SYN flood attack.

Managing IP Reputation policy settings

The FortiGuard IP Reputation service provides a database of known compromised or malicious client IP addresses. The database is updated periodically.

The IP Reputation configuration allows you to specify the action the system takes when an SLB virtual server receives traffic from a client with an IP address on the list. [Table 58](#) lists limitations for IP Reputation actions.

Table 58: IP Reputation actions

Action		Profile Limitations
Pass	IPv4 only	Not supported for RADIUS.
Deny	IPv4 only	Not supported for RADIUS.

Action		Profile Limitations
Redirect	IPv4 only	Not supported for RADIUS, FTP, TCP, UDP.
Send 403 Forbidden	IPv4 only	Not supported for RADIUS, FTP, TCP, UDP.

Note: IP Reputation is also not supported for Layer 4 virtual servers when the Packet Forwarding Mode is Direct Routing.

Basic Steps

1. Configure the connection to FortiGuard so the system can receive periodic IP Reputation Database updates. See [Configuring FortiGuard service settings](#).
2. Optionally, customize the actions you want to take when the system encounters a request from a source IP address that matches the list; and add exceptions. If a source IP address appears on the exceptions list, the system does not look it up on the IP Reputation list. See below.
3. Enable IP Reputation in the profiles you associate with virtual servers. See [Configuring Application profiles](#).

Before you begin:

- You must have Read-Write permission for Firewall settings.

To customize IP Reputation policy rules:

1. Go to Network Security > IP Reputation.
2. Make sure to select the **IP Reputation** tab, which displays all IP reputation policy configuration in FortiADC.
3. Click a policy or the corresponding Edit icon to open the IP Reputation editor.
4. Make the desired changes as described in [Table 59](#).
5. Click **Save**.

Table 59: IP Reputation policy configuration

Settings	Guidelines
Category	Depending the configuration on FortiGuard IP Reputation service, the IP reputation policy can be one of the following categories: <ul style="list-style-type: none"> • Botnet • Anonymous Proxy • Phishing • Spam • Other
Status	Enable or disable the category.

Settings	Guidelines
Action	<ul style="list-style-type: none"> • Pass • Deny • Redirect • Send 403 Forbidden <p>Note: Layer 4 and TCPS virtual servers do not support Redirect or Send 403 Forbidden. If you apply an IP Reputation configuration that uses these options to a Layer 4 or TCPS virtual server, FortiADC logs the action as Redirect or Send 403 Forbidden but in fact denies the traffic.</p>
Severity	<p>The severity to apply to the event. Severity is useful when you filter and sort logs:</p> <ul style="list-style-type: none"> • Low • Medium • High
Log	Enable or disable logging.

Configure IP reputation exception

To create an IP Reputation exception:

1. Go to Network Security > IP Reputation.
2. Click the **IP Reputation Exception** tab to add exceptions as described in [Table 60](#).
3. Click **Save**.

Table 60: IP Reputation exception

Settings	Guidelines
Status	Enable or disable the exception. You might have occasion to toggle the exception off and on.
Type	<ul style="list-style-type: none"> • IP/netmask: Select this option to block a specified IP address. • IP Range: Select this option to block a specified IP address range.
IP/Netmask	If IP/netmask is selected in the Type field above, specify a subnet using the address/mask notation.
Start IP / End IP	If IP Range is selected in the Type field above, specify the starting address and ending address of the IP range.

Using the Geo IP block list

The FortiGuard Geo IP service provides a database that maps IP addresses to countries, satellite providers, and anonymous proxies. The database is updated periodically.

The Geo IP block list is a policy that takes the action you specify when the virtual server receives requests from IP addresses in the blocked country's IP address space.

For Layer 4 virtual servers, FortiADC blocks access when the first TCP SYN packet arrives. For Layer 7 virtual servers, FortiADC blocks access after the handshake, allowing it to redirect the traffic if you have configured it to do so.

[Table 61](#) lists limitations for Geo IP block list actions.

Table 61: Geo IP block list actions

Action		Profile Limitations
Pass	IPv4 only	Not supported for HTTP Turbo, RADIUS.
Deny	IPv4 only	Not supported for HTTP Turbo, RADIUS.
Redirect	IPv4 only	Not supported for HTTP Turbo, RADIUS, FTP, TCP, TCPS, UDP.
Send 403 Forbidden	IPv4 only	Not supported for HTTP Turbo, RADIUS, FTP, TCP, TCPS, UDP.

Basic Steps

1. Configure the connection to FortiGuard so the system can receive periodic Geo IP Database updates. See [Configuring FortiGuard service settings](#).
2. Create rules to block traffic from locations.
3. Maintain a whitelist to allow traffic from specified subnets even if they belong to the address space blocked by the Geo IP block list.
4. Select the Geo IP block list and whitelist in the profiles you associate with virtual servers. See [Configuring Application profiles](#).

Before you begin:

- You must have Read-Write permission for Security settings.

To configure a Geo IP block list:

1. Go to Network Security > Geo IP Protection.
2. Click the **Geo IP Protection** tab.
3. Click **Add** to create a block list as described in [Table 62](#).
4. Click **Save**.

Table 62: Geo IP block list configuration

Settings	Guidelines
Name	Configuration name. Valid characters are A-Z, a-z, 0-9, _, and -. No spaces. After you initially save the configuration, you cannot edit the name.
Log	Enable/disable logging.
Action	<ul style="list-style-type: none"> • Pass—Allow the traffic. • Deny—Drop the traffic. • Redirect—Send a redirect. You specify the redirect URL on the profile configuration page. • Send 403 Forbidden—Send the HTTP Response code 403. <p>Note: Layer 4 and TCPS virtual servers do not support Redirect or Send 403 Forbidden. If you apply an Geo IP configuration that uses these options to a Layer 4 or TCPS virtual server, FortiADC logs the action as Redirect or Send 403 Forbidden, but in fact denies the traffic.</p>
Severity	The severity to apply to the event. Severity is useful when you filter and sort logs: <ul style="list-style-type: none"> • Low • Medium • High
Status	Enable or disable the Geo IP block list configuration.
Member	
Country	Select a geolocation object. The list includes countries as well as selections for anonymous proxies and satellite providers.

Using the Geo IP whitelist

To configure a Geo IP whitelist:

1. Go to Network Security > Geo IP Protection.
2. Click the **Whitelist** tab to create a whitelist as described in [Table 63](#).
3. Click **Save**.

Table 63: Geo IP whitelist configuration

Settings	Guidelines
Name	Configuration name. The name can be up to 35 characters long. Valid characters are A-Z, a-z, 0-9, _, and -. No space is allowed. After you initially save the configuration, you cannot edit the name.
Description	A string to describe the purpose of the configuration, to help you and other administrators more easily identify its use.
Status	Enable/disable the exception. You might have occasion to toggle the exception off and on.
Member	
Type	Select and configure either of the following: IP Subnet—Specify the IP address and CIDR-formatted subnet mask, separated by a forward slash (/), such as 192.0.2.0/24. Dotted quad formatted subnet masks are not accepted. IPv6 addresses are not supported. IP Range—Specify the Start IP and the End IP addresses of the IP range.
Description	Enter a brief description of the IP subnet or IP range, depending on which Type you choose. The description can be up to 1023 characters long. Valid characters are A-Z, a-z, 0-9, _, -, ., and :. No space is allowed.

Enabling denial of service protection

You can enable basic denial of service (DoS) prevention to combat [SYN floods](#). When enabled, FortiADC uses the SYN cookie method to track half-open connections. The system maintains a DoS mitigation table for each configured IPv4 virtual server. It times out half-open connections so that they do not deplete system resources.

Note: The DoS feature is supported for traffic to virtual servers only. However, it is not supported for IPv6 traffic or for Layer 4 virtual servers with the Direct Routing packet forwarding mode.

Before you begin:

- You must have Read-Write permission for Firewall settings.

To enable denial of service protection:

1. Go to Security > SYN Flood Prevention.
2. Enable the SYN Cookie feature.
3. Specify a maximum number of half open sockets. The default is 1 (10 connections). The valid range is 1 to 80,000.
4. Save the configuration.

Configuring a firewall policy

A firewall policy is a filter that allows or denies traffic based on a matching tuple: source address, destination address, and service. By default, firewall policy rules are stateful: if client-to-server traffic is allowed, the session is maintained in a state table, and the response traffic is allowed.

The FortiADC system evaluates firewall policies before other rules. It matches traffic against the firewall policy table, beginning with the first rule. If a rule matches, the specified action is taken. If the session is denied by a firewall policy rule, it is dropped. If the session is accepted, system processing continues.

By default, if firewall rules are not configured, the system does not perform firewall processing; all traffic is processed as if the system were a router, and traffic is forwarded according to routing and other system rules.

Note: You do not need to create firewall rules for routine management traffic associated with the management port or HA ports. The interface “allow access” option enables permitted protocols. The system automatically permits from-self traffic, such as health check traffic, and expected responses.

Before you begin:


- You must have a good understanding and knowledge of firewalls.
- You must have created the address configuration objects and service configuration objects that define the matching tuple in your firewall policy rules.
- You must have Read-Write permission for Firewall settings.

To configure a firewall:

1. Go to Network Security > Firewall.
2. Select [IPv4 Policy | IPv6 Policy].
3. Click **Add** to display the configuration editor.
4. Complete the configuration as described in [Table 64](#).
5. Save the configuration.
6. Reorder rules, as necessary.

Table 64: Firewall policy configuration

Settings	Guidelines
Default Action	Action when no rule matches or no rules are configured: <ul style="list-style-type: none"> • Deny—Drop the traffic. • Accept—Allow the traffic to pass the firewall.
Rule	
Name	Configuration name. Valid characters are A-Z, a-z, 0-9, _, and -. No spaces. After you initially save the configuration, you cannot edit the name.
Ingress Interface	Select the interface that receives traffic.

Settings	Guidelines
Egress Interface	Select the interface that forwards traffic.
Source	Select a source address object to use to form the matching tuple.
Destination	Select a destination address object to use to form the matching tuple.
Service	Select a service object to use to form the matching tuple.
Action	<ul style="list-style-type: none"> • Deny—Drop the traffic. • Accept—Allow the traffic to pass the firewall.
Reordering	
	After you have saved a rule, reorder rules as necessary. The rules table is consulted from top to bottom. The first rule that matches is applied and subsequent rules are not evaluated.

Configuring the firewall connection limit

The firewall connection limit policy allows or denies traffic based on a matching tuple: source address, destination address, and service; and connection count. The purpose is to detect anomalous connection requests.

The limit you specify can be based on the following counts:

- Count of concurrent sessions that match the tuple.
- Count of concurrent sessions from a single host that match the tuple.

The FortiADC system evaluates firewall connection limit policy rules before other rules. It matches traffic against the connection limit table, beginning with the first rule. If no rule matches, the connection is forwarded for further processing. If a rule matches, and the limit has not been reached, the connection is forwarded for further processing. If a rule matches and the limit has been reached, the connection is dropped.

By default, if firewall connection limit rules are not configured, the system does not perform connection limit policy processing. The firewall connection limit can be configured for non-SLB traffic and for Layer 7 SLB traffic, but not Layer 4 SLB traffic.

Note: The purpose of the firewall connection limit is distinct from the virtual server connection limit. The firewall connection limit setting is a security setting; the virtual server connection limit is a capacity setting.

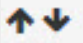
Before you begin:

- You must have a good understanding and knowledge of the capacity of your backend servers.
- You must have created the address configuration objects and service configuration objects that define the matching tuple in your connection limit rules.
- You must have Read-Write permission for Firewall settings.

To configure a firewall connection limit:

1. Click Network Security > [Firewall Connection Limit | IPv6 Connection Limit].
2. Click **Add** to display the configuration editor.
3. Complete the configuration as described in [Table 65](#).
4. Save the configuration.
5. Reorder rules, as necessary.

Table 65: Connection limit configuration

Settings	Guidelines
Name	Configuration name. Valid characters are A-Z, a-z, 0-9, _, and -. No spaces. After you initially save the configuration, you cannot edit the name.
Ingress Interface	Select the interface that receives traffic.
Egress Interface	Select the interface that forwards traffic.
Source	Select a source address object to use to form the matching tuple.
Destination	Select a destination address object to use to form the matching tuple.
Service	Select a service object to use to form the matching tuple.
Type	Specify whether the limit is per rule or per host.
Side	When the connection limit is per host, specify whether the connection counter gets incremented when the host IP address appears in: <ul style="list-style-type: none"> • Source—Only increment the counter if the host is the source address. • Destination—Only increment the counter if the host is the destination address. • Both—Increment the counter if the host is the source or destination address.
Limit	Maximum concurrent sessions. The default is 1,048,576.
Reordering	
	After you have saved a rule, reorder rules as necessary. The rules table is consulted from top to bottom. The first rule that matches is applied and subsequent rules are not evaluated.

Chapter 8: Web Application Firewall

You use web application firewall policies to scan HTTP requests and responses against known attack signatures and methods and filter matching traffic. This section includes the following topics:

- [Web application firewall basics](#)
- [Web application firewall configuration overview](#)
- [Configuring a WAF Profile](#)
- [Configuring a Web Attack Signature policy](#)
- [Configuring a URL Protection policy](#)
- [Configuring a HTTP Protocol Constraint policy](#)
- [Configuring an SQL/XSS Injection Detection policy](#)
- [Configuring WAF Exception objects](#)
- [Configuring a Bot Detection policy](#)

Web application firewall basics

A web application firewall (WAF) is a security policy enforcement point positioned between a client endpoint and a web application. The primary purpose is to prevent attacks against the web servers. A WAF is deployed separately from the web application so that the process overhead required to perform security scanning can be offloaded from the web server, and policies can be administered from one platform to many servers.

A WAF uses methods that complement perimeter security systems, such as the FortiGate next-generation firewall. The FortiADC WAF module applies a set of policies to HTTP scanpoints, which are parsed contexts of an HTTP transaction.

[Figure 47](#) illustrates the scanpoints. In the WAF policy configurations, you have options to enable rules to detect attacks at the request line, query string, filename, URI, request headers, request body, response code, or response body.

In particular:

- **Web Attack Signature policy**—The signature database includes signatures that can detect known attacks and exploits that can be found in 22 scanpoints. In your policy configuration, you choose classes of scanpoints to process: HTTP Headers, HTTP Request Body, and HTTP Response Body.
- **URL Protection policy**—This policy enables you to create rules that detect patterns in the URI or the file extension.
- **HTTP Protocol Constraint policy**—This policy enables you to create rules that restrict URI, header, and body length; HTTP method, or HTTP response code.
- **SQL/XSS Injection Detection policy**—This policy includes rules to detect SQL/XSS injection in the HTTP Request URI, HTTP Referer Header, HTTP Cookie Header, or HTTP Request Body.
- **Bot Detection**—This policy includes rules to detect Bots. A Bot is an application that runs automated tasks over the Internet. The WAF supports two methods for detecting bad Bots: signature detection and behavior detection. You can also use whitelists to exclude known trusted sources (good Bots) from detection.

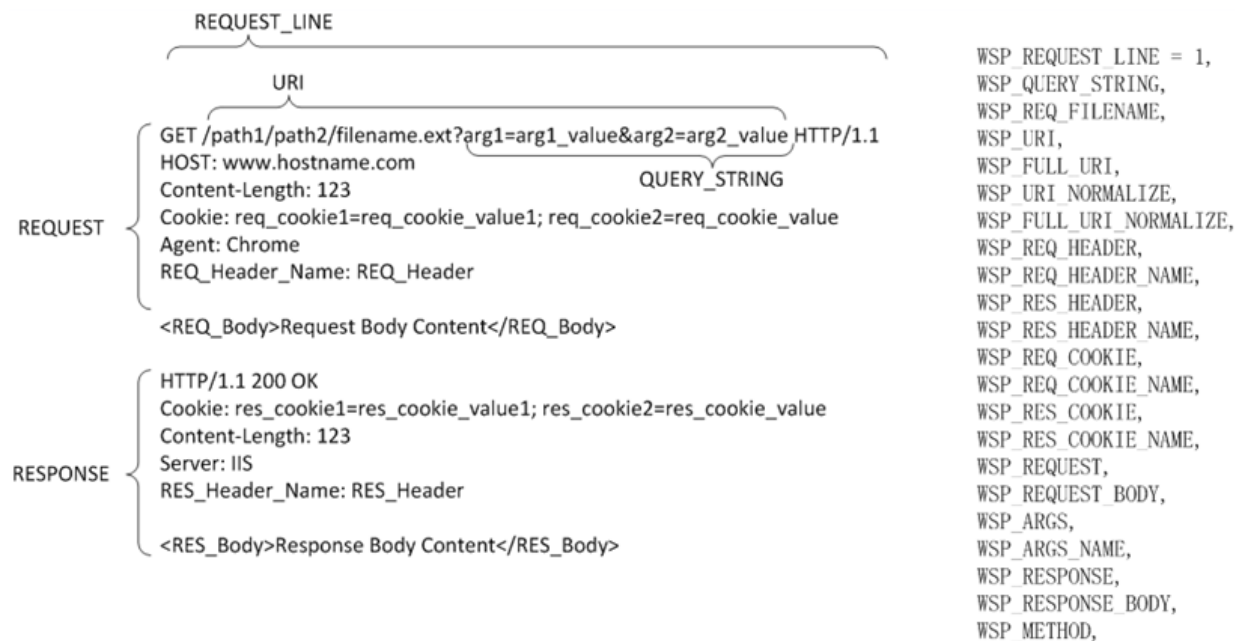
Policy rules are enforced (action taken) when scanning is completed at four checkpoints:

- HTTP Request Header
- HTTP Request Body

- HTTP Response Header
- HTTP Response Body

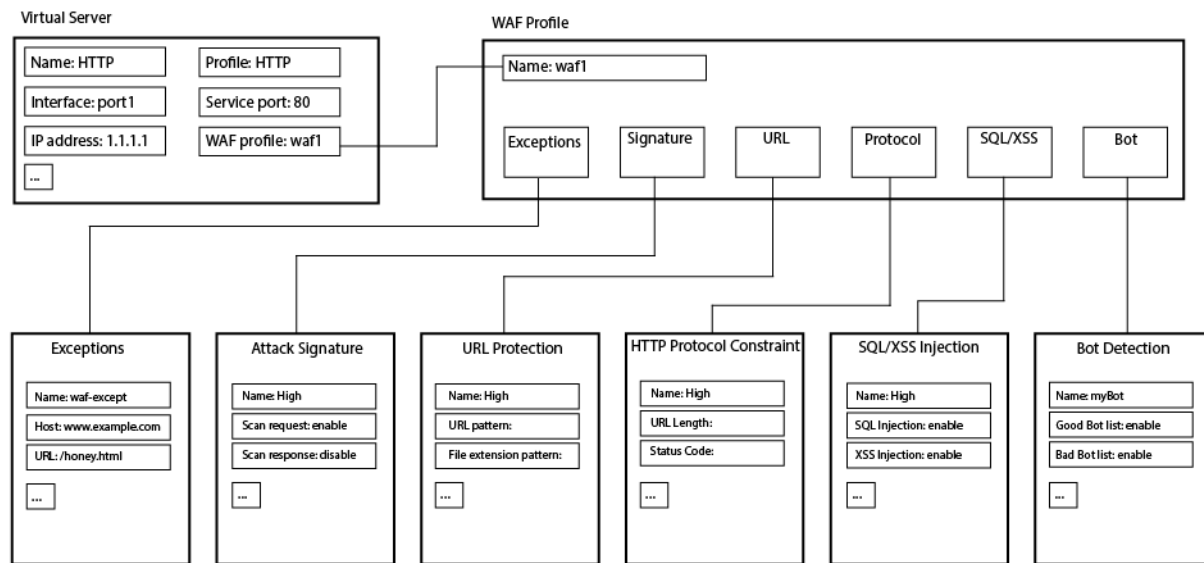
If the HTTP Request Header violates a rule, and the action is Deny, the attempted session is dropped and scanning for the transaction stops. If the action is Alert, the event is logged and rules processing continues.

Figure 47: HTTP scanpoints



Web application firewall configuration overview

Figure 48 shows the relationship between WAF configuration elements. A WAF profile comprises a Web Attack Signature policy, URL Protection policy, HTTP Protocol Constraint policy, SQL/XSS Injection Detection, and Bot Detection policy. The profile is applied to a load balancing virtual server, so all traffic routed to the virtual server is subject to the WAF rules. WAF profiles can be applied to HTTP and HTTPS virtual servers but not HTTP Turbo virtual servers.

Figure 48: WAF configuration overview

Predefined configuration elements

The FortiADC WAF includes many predefined configuration elements to help you get started. It includes predefined WAF profiles, predefined Web Attack Signature policies, predefined HTTP Protocol Constraint policies, and predefined SQL/XSS Injection Detection policies.

Severity

The severity ratings for predefined Web Attack Signatures and the default severity rating for feature options like SQL/XSS Injection Detection are based on the Open Web Application Security Project (OWASP) [Risk Rating Methodology](#). In order to harmonize the significance of severity levels in logs, we recommend you use this methodology to assign severity for any custom elements you create.

Exceptions

You can create exceptions so that traffic to specific hosts or URL patterns is not subject to processing by WAF rules. Exception lists are processed before traffic is inspected. If an exception applies, the traffic bypasses the WAF module.

Basic Steps

1. Create configuration objects that define the exception.
2. Add the exception to a WAF profile configuration or WAF rule configuration.

Configuring a WAF Profile

A WAF profile references the WAF policies that are to be enforced.

[Table 66](#) describes the predefined profiles. In many cases, you can use predefined profiles to get started.

Table 66: Predefined WAF profiles

Predefined Profiles	Description
High-Level-Security	<ul style="list-style-type: none"> • Web Attack Signature policy: High-Level-Security • HTTP Protocol Constraints policy: High-Level-Security • SQL/XSS Injection Detection policy: High-Level-Security
Medium-Level-Security	<ul style="list-style-type: none"> • Web Attack Signature policy: Medium-Level-Security • HTTP Protocol Constraints policy: Medium-Level-Security • SQL/XSS Injection Detection policy: Medium-Level-Security
Alert-Only	<ul style="list-style-type: none"> • Web Attack Signature policy: Alert-Only • HTTP Protocol Constraints policy: Alert-Only • SQL/XSS Injection Detection policy: Alert-Only

If desired, you can create user-defined profiles. The maximum number of profiles per VDOM is 255.

Before you begin:

- You can use predefined WAF profiles, create profiles based on predefined feature options, or create profiles based on user-defined configuration objects. If you want to add user-defined configuration objects, you must create them before using this procedure to add them to a WAF profile.
- You must have Read-Write permission for Security settings.

After you have created a WAF profile, you can specify it in a virtual server configuration.

To configure a WAF Profile:

1. Go to Security > Web Application Firewall.
2. Click the **WAF Profile** tab.
3. Click Add to display the configuration editor.
4. Complete the configuration as described in [Table 67](#).
5. Save the configuration.

Table 67: WAF Profile configuration

Settings	Guidelines
Name	Configuration name. Valid characters are A-Z, a-z, 0-9, _, and -. No spaces. After you initially save the configuration, you cannot edit the name.
Description	A string to describe the purpose of the configuration, to help you and other administrators more easily identify its use.
Web Attack Signature	Select a predefined or user-defined Web Attack Signature configuration object.

Settings	Guidelines
URL Protection	Select a user-defined URL Protection configuration object.
HTTP Protocol Constraint	Select a predefined or user-defined HTTP Protocol Constraint configuration object.
SQL/XSS Injection Detection	Select a predefined or user-defined SQL/XSS Injection Detection configuration object.
Exception Name	Select a user-defined exception configuration object. Exceptions identify specific hosts or URL patterns that are not subject to processing by this rule.
Bot Detection	Select a user-defined Bot Detection configuration object.

Configuring a Web Attack Signature policy

The FortiGuard Web Attack Signature service provides a database of attack signatures that is updated periodically to protect against new kinds of attacks. [Table 70](#) summarizes the categories of threats that are detected by the signatures. The categories are reported in logs.

In the Web Attack Signature policy configuration, you can enable/disable the class of scanpoints and the action when traffic matches signatures.

There are three classes of scanpoints:

- HTTP Header—Scans traffic against HTTP header signatures. If you enable a policy at all, you are enabling HTTP header scanning.
- HTTP Request Body—Scans traffic against HTTP request body signatures.
- HTTP Response Body—Scans traffic against HTTP response body signatures.

Header scanning is always a good practice, so enabling a policy always enables header scanning. Body scanning impacts performance, so you have the option of disabling body scanning if system utilization or latency become an issue.

You can specify separate actions for three levels of event severity:

- High—We recommend you deny traffic for high severity events.
- Medium—We recommend you deny or alert, according to your preference. To be strict, deny; otherwise, alert.
- Low—We recommend you allow the traffic and log an alert for low severity events.

[Table 68](#) describes the predefined policies. You can select the predefined policies in your WAF profiles, or you can create policies that enable a different set of scan classes or a different action. In this release, you cannot exclude individual signatures or create custom signatures. You can enable or disable the scan classes.

Table 68: Web Attack Signature predefined policies

Policy	Status	Action
High-Level-Security	Scan HTTP header—Enabled.	High Severity Action—Deny.
	Scan HTTP Request Body—Enabled.	Medium Severity Action—Deny.
	Scan HTTP Response Body—Disabled.	Low Severity Action—Alert.
Medium-Level-Security	Scan HTTP header—Enabled.	High Severity Action—Deny.
	Scan HTTP Request Body—Enabled.	Medium Severity Action—Alert.
	Scan HTTP Response Body—Disabled.	Low Severity Action—Alert.
Alert-Only	Scan HTTP header—Enabled.	High Severity Action—Alert.
	Scan HTTP Request Body—Disabled.	Medium Severity Action—Alert.
	Scan HTTP Response Body—Disabled.	Low Severity Action—Alert.

Basic Steps

1. Configure the connection to FortiGuard so that the system can receive periodic WAF Signature Database updates. See [Configuring FortiGuard service settings](#).
2. Optionally, if you do not want to use the predefined policies, configure Web Attack Signature policies. See below.
3. When configuring the WAF profile, select a policy that you associate with virtual servers . See [Configuring a Web Attack Signature policy](#).

Before you begin:

- You must have read-write permission for security settings.

To configure a Web Attack Signature policy:

1. Go to Web Application Firewall.
2. Click the **Web Attack Signature** tab.
3. Click **Add** to display the configuration editor.
4. Complete the configuration as described in [Table 69](#).
5. Save the configuration.

Table 69: Web Attack Signature configuration

Settings	Guidelines
Category	This dialog provides tools for configuring a Web attack signature policy.

Settings	Guidelines
Name	<p>Specify a unique name for the Web attack signature policy and click Save. Valid characters are A-Z, a-z, 0-9, _, and -. No space is allowed between characters.</p> <p>Note: Once saved, the policy name cannot be changed.</p>
Category	<p>This section lists the (main) categories of Web attack signatures within the system. Do the following to include the desired categories of Web attack signature in the policy:</p> <ol style="list-style-type: none"> 1. In the Name column, identify the categories of Web attack signatures of interest. 2. In the Status column, select (check mark) the categories you like to include in the policy. 3. In the Action column, select the action you want to apply to the categories that you select. 4. Double-click the name of a category to view its sub-categories. See Sub-category below.
Sub-category	<p>This section lists the sub-categories of a (main) category of Web attack signature that you have opened (double-clicked) from above. Do the following to enable any of the sub-categories of interest:</p> <ol style="list-style-type: none"> 1. In the Name column, identify the sub-categories of interest. 2. In the Status column, select (check mark) the sub-categories you like to include in the policy.
Signature	<p>This dialog provides tools for searching through and filtering Web attack signatures available within the system.</p>
Search	<p>Use the following options to search for Web attack signatures to display:</p> <ul style="list-style-type: none"> • Description—Enter a descriptive text string and click Search. • ID—Enter a Web attack signature ID and click Search. • CVE Number—Enter a CVE number related to a Web attack signature and click Search. • Clear Search—Click this button to empty all search fields. <p>Note: Web attack signatures that match your search criterion show up in the Signature section below the moment you click the corresponding Search button.</p>

Settings	Guidelines
Filters	<p>Use any or a combination of the following filters to filter the Web attack signatures to be displayed in the Signature section below:</p> <ul style="list-style-type: none"> • Category—Click the down arrow and select a (main) category of Web attack signatures from the drop-down menu. • Sub-category—Click the down arrow and select a sub-category of the category of Web attack signatures that you have selected. • Status—Click the down arrow and select either (Enable or Disable) from the drop-down menu. • Severity—Click the down arrow and select High, Medium, or Low from the drop-down menu. • With Exception—Click the down arrow and select either (Yes or No) from the drop-down menu. • Clear All—Click this button to clear the existing filters. Note: You can also remove a specific filter by clicking the corresponding x mark.
Signature	<p>This section displays all Web attack signatures that match your search and filter criteria, showing the following information for each Web attack signature:</p> <ul style="list-style-type: none"> • ID • Status • Name • Severity • Target Application • Exception Name
Signature Detail	<p>This section shows detailed information about the Web attack signature that you've highlighted (clicked) in the Signature section above.</p>
Detail	<p>This tab shows the following information about the selected signature:</p> <ul style="list-style-type: none"> • Signature ID • Category • Sub-category • Severity • Target Application • Description
Edit Signature	<p>This tab provides tools for editing a selected Web attack signature. It contains the following fields:</p> <ul style="list-style-type: none"> • Signature ID—(Read only) Shows the ID of the selected signature. • Status—Click to enable or disable the signature. • Exception Name—Click the down arrow and select an exception from the drop-down menu.

Table 70 summarizes the categories of threats that are detected by the signatures.

Table 70: Web Attack Signature categories and subcategories

Category (ID)	Subcategory (ID)
Cross Site Scripting (1)	
SQL Injection (2)	
Generic Attacks (3)	OS Command Injection (1)
	Coldfusion Injection (2)
	LDAP Injection (3)
	Command Injection (4)
	Session Fixation (5)
	File Injection (6)
	PHP Injection (7)
	SSI Injection (8)
	UPDF XSS (9)
	Email Injection (10)
	HTTP Response Splitting (11)
	RFI Injection (12)
Trojans (4)	

Category (ID)	Subcategory (ID)
Information Disclosure (5)	Zope Information Leakage (13)
	CF Information Leakage (14)
	PHP Information Leakage (15)
	ISA Server Existence Revealed (16)
	Microsoft Office Document Properties Leakage (17)
	CF Source Code Leakage (18)
	IIS Information Leakage (19)
	Weblogic information leakage (20)
	Generic Filename and Directory leakage (21)
	ASP/JSP Source Code Leakage (22)
	PHP Source Code Leakage (23)
	SQL Error leakage (24)
	HTTP Header Leakage (25)
	WordPress Leakage (26)

Category (ID)	Subcategory (ID)
Known Exploits (6)	Oracle 9i (27)
	Coppermine Photo Gallery (28)
	Netscape Enterprise Server (29)
	Cisco IOS HTTP Service (30)
	Microsoft SQL Server (31)
	HP OpenView Network Node Manager (32)
	Best Software SalesLogix (33)
	IBM Lotus Domino Web Server (34)
	Microsoft IIS (35)
	Microsoft Windows Media Services (36)
	Dave Carrigan Auth_LDAP (37)
	427BB (38)
	RaXnet Cacti Graph (39)
	CHETCPASSWD (40)
	SAP (41)
Credit Card Detection (7)	
Bad Robot (8)	

Configuring a URL Protection policy

URL protection policies can filter HTTP requests that match specific character strings and file extensions.

Before you begin:

- You must have Read-Write permission for Security settings.

After you have configured URL protection policies, you can select them in WAF profiles.

To configure a URL Protection policy:

1. Go to Security > Web Application Firewall.
2. Click the **URL Protection** tab.
3. Click **Add** to display the configuration editor.

4. Complete the configuration as described in [Table 71](#).
5. Save the configuration.

Table 71: URL Protection configuration

Settings	Guidelines
Name	Configuration name. Valid characters are A-Z, a-z, 0-9, _, and -. No spaces. After you initially save the configuration, you cannot edit the name.
URL Access Rule	
Full URL Pattern	Matching string. Regular expressions are supported.
Action	<ul style="list-style-type: none"> Alert—Allow the traffic and log the event. Deny—Drop the traffic, send a 403 Forbidden to the client, and log the event. The default is alert.
Severity	<ul style="list-style-type: none"> High—Log as high severity events. Medium—Log as a medium severity events. Low—Log as low severity events. The default is low.
Exception Name	Select an exception configuration object. Exceptions identify specific hosts or URL patterns that are not subject to processing by this rule.
File Extension Rule	
File Extension Pattern	Matching string. Regular expressions are supported.
Action	<ul style="list-style-type: none"> Alert—Allow the traffic and log the event. Deny—Drop the traffic, send a 403 Forbidden to the client, and log the event. The default is alert.
Severity	<ul style="list-style-type: none"> High—Log as high severity events. Medium—Log as a medium severity events. Low—Log as low severity events. The default is low.
Exception Name	Select an exception configuration object. Exceptions identify specific hosts or URL patterns that are not subject to processing by this rule.

Configuring an HTTP Protocol Constraint policy

The HTTP Protocol Constraint policy includes the following rules:

- HTTP request parameters—Limit the length of URIs, headers, and body to prevent several types of attacks, such as buffer overflow and denial of service.
- HTTP request methods—Restrict [HTTP methods](#) allowed in HTTP requests. For example, do not allow the PUT method in HTTP requests to prevent attackers from uploading malicious files.
- HTTP response codes—Drop response traffic containing [HTTP response codes](#) that might contain information attackers can use to craft attacks. For example, some HTTP response codes include fingerprint data like web server version, database version, OS, and so on.

[Table 72](#) describes the predefined policies.

Table 72: Predefined HTTP protocol constraint policies

Predefined Rules	Description
High-Level-Security	Protocol constraints enabled with default values. Action is set to deny. Severity is set to high.
Medium-Level-Security	Protocol constraints enabled with default values. Action is set to alert. Severity is set to medium.
Alert-Only	Protocol constraints enabled with default values. Action is set to alert. Severity is set to low.

If desired, you can create user-defined rules to filter traffic with invalid HTTP request methods or drop packets with the specified server response codes.

Before you begin:

- You should have a sense of legitimate URI lengths and HTTP request methods for the destination resources.
- You should know whether your servers include application fingerprint information in HTTP response codes.
- You must have Read-Write permission for Security settings.

To configure an HTTP Protocol Constraint policy:

1. Go to Security > Web Application Firewall.
2. Click the **HTTP Protocol Constraint** tab.
3. Click **Add** to display the configuration editor.
4. Complete the configuration as described in [Table 73](#).
5. Save the configuration.

Table 73: HTTP Protocol Constraint configuration

Settings	Guidelines
Name	<p>Enter a unique HTTP protocol constraint policy name. Valid characters are A-Z, a-z, 0-9, _, and -. No space is allowed.</p> <p>Note: Once saved, the name of an HTTP protocol constraint policy cannot be changed.</p>
Request Parameters	

Settings	Guidelines
Maximum URI Length	Maximum characters in an HTTP request URI. The default is 2048. The valid range is 1-8192.
Illegal Host Name	Enable/disable hostname checks. A domain name must consist of only the ASCII alphabetic and numeric characters, plus the hyphen. The hostname is checked against the set of characters allowed by the RFC 2616. Disallowed characters, such as non-printable ASCII characters or other special characters (for example, '<', '>', and the like), are a symptom of an attack.
Illegal HTTP Version	Enable/disable the HTTP version check. Well-formed requests include the version of the protocol used by the client, in the form of HTTP/v where v is replaced by the actual version number (one of 0.9, 1.0, 1.1). Malformed requests are a sign of traffic that was not sent from a normal browser and are a symptom of an attack.
Illegal HTTP Multipart	Enable/Disable the HTTP body multipart check. If the content-type is multipart media type, the HTTP body must contain one or more body parts, each preceded by a boundary delimiter line and the last one followed by a closing boundary delimiter line. After its boundary delimiter line, each body part then consists of a header area, a blank line, and a body area. Malformed HTTP requests are a sign of traffic that was not sent from a normal browser and are a symptom of an attack.
Maximum Cookie Number In Request	Maximum number of cookie headers in an HTTP request. The default is 16. The valid range is 1-32.
Maximum Header Number In Request	Maximum number of headers in an HTTP request. The default is 50. Requests with more headers are a symptom of a buffer overflow attack or an attempt to evade detection mechanisms. The valid configuration range is 1-100.
Maximum Request Header Name Length	Maximum characters in an HTTP request header name. The default is 1024. The valid range is 1-8192.
Maximum Request Header Value Length	Maximum characters in an HTTP request header value. The default is 4096. Longer headers might be a symptom of a buffer overflow attack. The valid configuration range is 1-8192.
Maximum URL Parameter Name Length	Maximum characters in a URL parameter name. The default is 1024. The valid range is 1-2048.
Maximum URL Parameter Value Length	Maximum characters in a URL parameter value. The default is 4096. The valid range is 1-8192.
Maximum Request Header Length	Maximum length of the HTTP request header. The default is 8192. The valid range is 1-16384.

Settings	Guidelines
Maximum Request Body Length	Maximum length of the HTTP body. The default is 67108864. The valid range is 1-67108864.
Request Method Rule	
Method	<p>Select one or more methods to match in the HTTP request line:</p> <ul style="list-style-type: none"> • CONNECT • DELETE • GET • HEAD • OPTIONS • POST • PUT • TRACE • Others <p>Note: The first 8 methods are described in RFC 2616. The group Others contains not commonly used HTTP methods defined by Web Distributed Authoring and Version (WebDAV) extensions.</p>
Action	<ul style="list-style-type: none"> • Alert—Allow the traffic and log the event. • Deny—Drop the traffic, send a 403 Forbidden to the client, and log the event. <p>The default is alert.</p>
Severity	<ul style="list-style-type: none"> • High—Log as high severity events. • Medium—Log as a medium severity events. • Low—Log as low severity events. <p>The default is low.</p>
Exception	Select an exception configuration object. Exceptions identify specific hosts or URL patterns that are not subject to processing by this rule.
Response Code Rule	
Minimum Status Code / Maximum Status Code	Start/end of a range of status codes to match. You can specify codes 400 to 599.
Action	<ul style="list-style-type: none"> • Alert—Allow the traffic and log the event. • Deny—Drop the traffic, send a 403 Forbidden to the client, and log the event. <p>The default is alert.</p>

Settings	Guidelines
Severity	<ul style="list-style-type: none"> • High—Log as high severity events. • Medium—Log as a medium severity events. • Low—Log as low severity events. <p>The default is low.</p>
Exception	Select an exception configuration object. Exceptions identify specific hosts or URL patterns that are not subject to processing by this rule.

Configuring an SQL/XSS Injection Detection policy

SQL/XSS Injection Detection policies detect [SQL injection](#) and [cross-site scripting \(XSS\)](#) attacks. Injection occurs when user-supplied data is sent to an interpreter as part of a command or query. In an SQL injection attack, attackers craft HTTP requests that cause SQL queries to be executed directly against the web application's database. XSS injection attacks cause a web browser to execute a client-side script.

In contrast to signature-based detection, the WAF SQL and XSS injection detector module detects SQL and XSS injection through lexical analysis, which is a complementary method and is faster.

The policy enables/disables scanpoints, the action when traffic matches signatures, and the event severity.

You can enable detection in the following scanpoints:

- SQL Injection: URI—Analyzes content in the URI.
- SQL Injection: Referer—Analyzes content in the HTTP Referer header.
- SQL Injection: Cookie—Analyzes content in the HTTP Cookie header.
- SQL Injection: Body—Analyzes content in the HTTP request body.
- XSS Injection: URI—Analyzes content in the URI.
- XSS Detection: Referer—Analyzes content in the HTTP Referer header.
- XSS Detection: Cookie—Analyzes content in the HTTP Cookie header.
- XSS Detection: Body—Analyzes content in the HTTP request body.

Header scanning is recommended. Body scanning impacts performance, so you have the option of disabling body scanning if system utilization or latency become an issue.

[Table 74](#) describes the predefined policies.

Table 74: Predefined SQL injection and XSS detection policies

Predefined Rules	SQL Injection			XSS		
	Detection	Action	Severity	Detection	Action	Severity
High-Level-Security	All except Body SQL Injection Detection	Deny	High	All except Body XSS Injection Detection	Deny	High

Predefined Rules	SQL Injection			XSS		
	Detection	Action	Severity	Detection	Action	Severity
Medium-Level-Security	Only SQL URI SQL Injection Detection	Deny	High	None	Alert	Low
Alert-Only	Only SQL URI SQL Injection Detection	Alert	High	None	Alert	Low

If desired, you can create user-defined policies.

Before you begin:

- You must have Read-Write permission for Security settings.

After you have created an SQL injection/XSS policy, you can specify it in a WAF profile configuration.

To configure an SQL/XSS Injection Detection policy:

- Go to Security > Web Application Firewall.
- Click the **SQL/XSS Injection Detection** tab.
- Click **Add** to display the configuration editor.
- Complete the configuration as described in [Table 75](#).
- Save the configuration.

Table 75: SQL/XSS Injection Detection configuration

Settings	Guidelines
Name	Configuration name. Valid characters are A-Z, a-z, 0-9, _, and -. No spaces. After you initially save the configuration, you cannot edit the name.
SQL	
SQL Injection Detection	Enable/disable SQL injection detection.
URI Detection	Enable/disable detection in the HTTP request.
Referer Detection	Enable/disable detection in the Referer header.
Cookie Detection	Enable/disable detection in the Cookie header.

Settings	Guidelines
Body Detection	Enable/disable detection in the HTTP Body message.
Action	<ul style="list-style-type: none"> • Alert—Allow the traffic and log the event. • Deny—Drop the traffic, send a 403 Forbidden to the client, and log the event. <p>The default is alert, but we recommend you deny SQL Injection.</p>
Severity	<ul style="list-style-type: none"> • High—Log as high severity events. • Medium—Log as a medium severity events. • Low—Log as low severity events. <p>The default is low, but we recommend you rate this high or medium.</p>
SQL Exception Name	Select an exception configuration object. Exceptions identify specific hosts or URL patterns that are not subject to processing by this rule.
XSS	
XSS Injection Detection	Enable/disable XSS injection detection.
URI Detection	Enable/disable detection in the HTTP request.
Referer Detection	Enable/disable detection in the Referer header.
Cookie Detection	Enable/disable detection in the Cookie header.
Body Detection	Enable/disable detection in the HTTP Body message.
Action	<ul style="list-style-type: none"> • Alert—Allow the traffic and log the event. • Deny—Drop the traffic, send a 403 Forbidden to the client, and log the event. <p>The default is alert, but we recommend you deny XSS Injection.</p>
Severity	<ul style="list-style-type: none"> • High—Log matches as high severity events. • Medium—Log matches as a medium severity events. • Low—Log matches as low severity events. <p>The default is low, but we recommend you rate this high or medium.</p>
XSS Exception Name	Select an exception configuration object. Exceptions identify specific hosts or URL patterns that are not subject to processing by this rule.

Configuring WAF Exception objects

Exceptions identify specific hosts or URL patterns that are not subject to processing by WAF rules.

Before you begin:

- You must have Read-Write permission for Security settings.

After you have created an exception object, you can specify it in WAF profiles and individual WAF feature rules.

To configure an exception object:

1. Go to Security > Web Application Firewall.
2. Click the **Exceptions** tab.
3. Click **Add** to display the configuration editor.
4. Complete the configuration as described in [Table 76](#).
5. Save the configuration.

Table 76: WAF Exception objects

Settings	Guidelines
Name	Configuration name. Valid characters are A-Z, a-z, 0-9, _, and -. No spaces. After you initially save the configuration, you cannot edit the name.
Exception Host Status	Enable/disable setting exceptions by host pattern.
Exception Host	Matching string. Regular expressions are supported. For example, you can specify <code>www.example.com</code> , <code>*.example.com</code> , or <code>www.example.*</code> to match a literal host pattern or a wildcard host pattern.
Exception URL	Matching string. Must begin with a URL path separator (/). Regular expressions are supported. For example, you can specify pathnames and files with expressions like <code>\admin</code> , <code>.*\data\1.html</code> , or <code>\data.*</code> .

Configuring a Bot Detection policy

Bot detection policies use signatures and source behavior tracking to detect client traffic likely to be generated by robots instead of genuine clients. Some bots, such as search engine crawlers, are "good bots" that perform search indexing tasks that can result in more legitimate users being directed to your site. You enable a whitelist to permit those. "Bad bots" are known to send traffic that has a negative impact on site availability and integrity, such as DDoS attacks or content scrapping. You want to block these.

To get started, you can use predefined whitelists (known good bots) and blacklists (known bad bots). You can also specify a rate limit threshold of HTTP requests/second for sources not matched to either whitelist or blacklist. The rate limit threshold can be useful in detecting "unknown bots".

In the event of false positives, you can use the user-specified whitelist table to fine-tune detection.

Before you begin:

- You must configure the connection to FortiGuard so the system can receive periodic WAF Signature Database updates, including "good bot" and "bad bot" signatures and lists. See [Configuring FortiGuard service settings](#).
- You must have Read-Write permission for Security settings.

After you have configured Bot Detection policies, you can select them in WAF profiles.

To configure a Bot Detection policy:

1. Go to Security > Web Application Firewall.
2. Click the **Bot Detection** tab.
3. Click **Add** to display the configuration editor.
4. Complete the configuration as described in [Table 77](#).
5. Save the configuration.

Table 77: Bot Detection configuration

Settings	Guidelines
Name	Configuration name. Valid characters are A-Z, a-z, 0-9, _, and -. No spaces. After you initially save the configuration, you cannot edit the name.
Status	Enable/disable Bot detection.
Search Engine Status	Enable/disable the predefined search engine spider whitelist. The list is included in WAF signature updates from FortiGuard.
Bad Robot Status	Enable/disable the predefined bad robot blacklist. The list is included in WAF signature updates from FortiGuard.
HTTP Request Rate	Specify a threshold (HTTP requests/second/source) to trigger the action. Bots send HTTP request traffic at extraordinarily high rates. The source is tracked by source IP address and User-Agent. The default is 0 (off). The valid range is 0-100,000,000 requests per second.
Action	<ul style="list-style-type: none"> • Alert—Allow the traffic and log the event. • Deny—Drop the traffic, send a 403 Forbidden to the client, and log the event. The default is alert.
Severity	<ul style="list-style-type: none"> • High—Log as high severity events. • Medium—Log as a medium severity events. • Low—Log as low severity events. The default is low.
Block Period	The default is 3600 seconds. The valid range is 1-3600. The maximum size of the block IP address table is 100,000 entries. If the table is full, the earliest entry will be deleted.

Settings	Guidelines
Whitelist	
IPv4/Netmask	Matching subnet (CIDR format).
URL Pattern	Matching string. Regular expressions are supported.
URL Parameter Name	Matching string. Regular expressions are supported.
Cookie Name	Matching string. Regular expressions are supported.
User Agent	Matching string. Regular expressions are supported.

Chapter 9: Authentication Management

This chapter includes the following topics:

- [Managing administrator users](#)
- [Using the local authentication server](#)
- [Using a RADIUS authentication server](#)
- [Using an LDAP authentication server](#)
- [Using Kerberos authentication relay](#)
- ["Using HTTP Basic SSO" on page 252](#)
- [Configure SAML authentication](#)
- [Configuring user groups](#)

Configuring user groups

User groups are authorized by the virtual server authentication policy. The user group configuration references the authentication servers that contain valid user credentials.

Suggested steps:

1. Configure LDAP and RADIUS servers, if applicable.
2. Configure local users.
3. Configure user groups (reference servers and local users).
4. Configure an authentication policy (reference the user group).
5. Configure the virtual server (reference the authentication policy).

Before you begin:

- You must have created configuration objects for any LDAP and RADIUS servers you want to use, and you must have created user accounts for local users.
- You must have read-write permission for System settings.

After you have created user groups, you can specify them in the server load balancing authentication policy configuration.

To configure a user group:

1. Go to Authentication Management > User Group.
2. Click **Add** to display the configuration editor.
3. Complete the configuration as described in [Table 78](#).
4. Save the configuration.

Table 78: User group configuration

Settings	Guidelines
User Group	Use this page to configure a user group.
Name	<p>Specify a unique name for the user group. Valid characters are A-Z, a-z, 0-9, _, and -. No spaces.</p> <p>After you initially save the configuration, you cannot edit the name.</p>
User Cache	Enable to cache the credentials for the remote users (LDAP, RADIUS) once they are authorized.
Cache Timeout	Timeout for cached user credentials. The default is 300 seconds. The valid range is 1-86,400 seconds.
Authentication Timeout	Timeout for query sent from FortiADC to a remote authentication server. The default is 2,000 milliseconds. The valid range is 1-60,000 milliseconds.
Authentication Log	<p>Specify one of the following logging options for authentication events:</p> <ul style="list-style-type: none"> • No logging • Log failed attempts • Log successful attempts • Log all (both failed and successful attempts)
Client Authentication Method	<ul style="list-style-type: none"> • HTML Form • HTTP
Group Type	<ul style="list-style-type: none"> • Normal—Default. No action is needed. • SSO—Select to enable single sign-on (SSO) and then populate the fields below.
Authentication Relay	Select an authentication relay profile.
Authentication Session Timeout	Specify the authentication session timeout. Valid values range from 1 to 180 (minutes). The default is 3 (minutes).
SSO Support	Disabled by default. When enabled, you must specify the SSO domain. See below.
SSO Domain	<p>Specify the SSO domain.</p> <p>Note: This field becomes available only when SSO Support is enabled. See above.</p>

Settings	Guidelines
Log-off URL	Specify the log-off URL.
Member	Use this page to add members to a user group,
Type	Select one of the following: <ul style="list-style-type: none"> • Local • LDAP • RADIUS
Local User	If Type is set to be Local (see above), select a local user (account name).
LDAP Server	If Type is set to be LDAP Server (see above), select an LDAP server.
RADIUS Server	If Type is set to be RADIUS Server (see above), select a RADIUS server.

Using the local authentication server

You can use a local authentication server to authenticate destination server user logins.

Note: The local authentication server does not have user-initiated password management features, so it does not easily scale to large groups of users. For large deployments, we recommend using RADIUS or LDAP and providing instructions for users on how to reset, recover, or change their passwords.

Basic steps:

1. Add user accounts to the local authentication server.
2. Select the local authentication server configuration and username when you create user groups.

Before you begin:

- You must have read-write permission for system settings.

To use a local authentication server:

1. Go to Authentication Management > Local User.
2. Click **Add** to display the configuration editor.
3. Complete the configuration as described in [Table 79](#).
4. Save the configuration.

Table 79: Local authentication server configuration

Settings	Guidelines
Name	Specify a unique name for the user account, such as <code>user1</code> or <code>user1@example.com</code> . Note: A user name can contain up to 35 characters. Use no space or special characters except the 'at' symbol (@) or dot (.). Once saved, the user name cannot be changed.
Password	Create a password. The stored password will be encrypted.

Using an LDAP authentication server


Lightweight Directory Access Protocol (LDAP) is an application protocol for accessing and maintaining distributed directory information services over a network. When using LDAP, authentication clients may send “Bind” messages to servers for authentication. Depending on the circumstances, clients may send different kinds of “Bind” messages.

LDAP bind messages

In a server load-balancing client authentication or admin authentication scenario, FortiADC sends a binding request to the LDAP server for client authentication. Once a client is successfully authenticated, he or she can then access the LDAP server based on his or her privileges. There are three bind types: simple, anonymous, and regular.

Simple bind

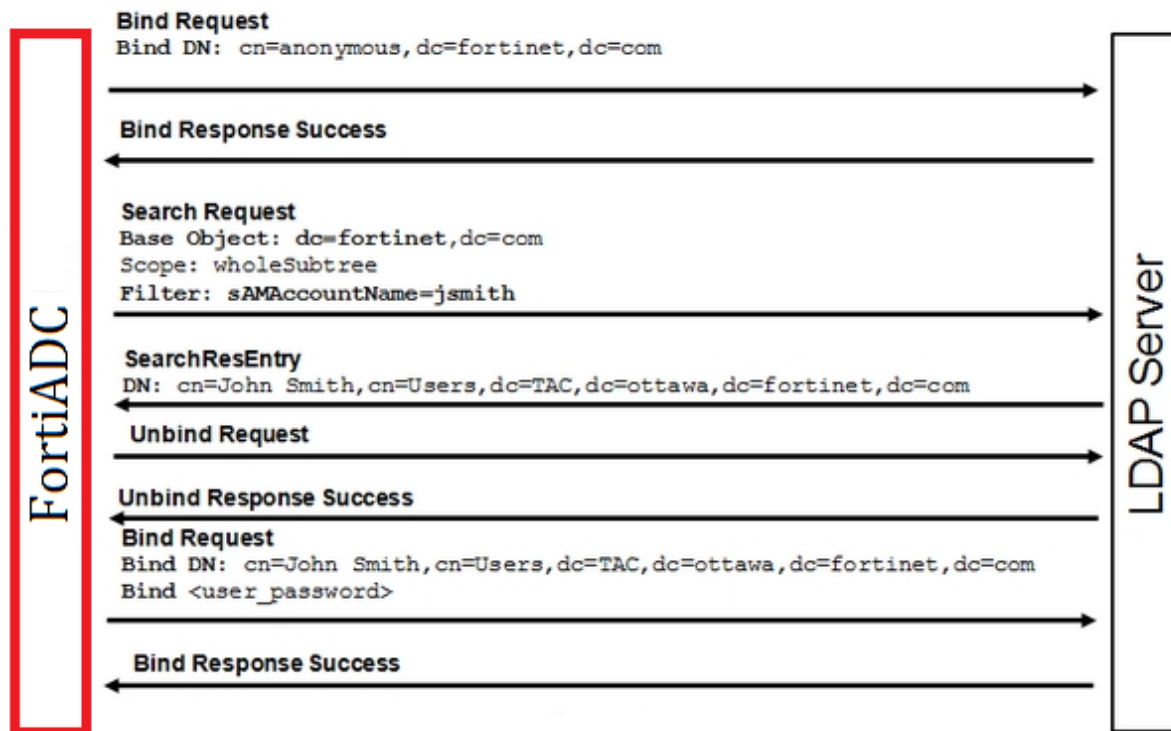
Simple bind means binding with a client's full name. All clients must be located in the same branch specified with the DN.



Anonymous bind

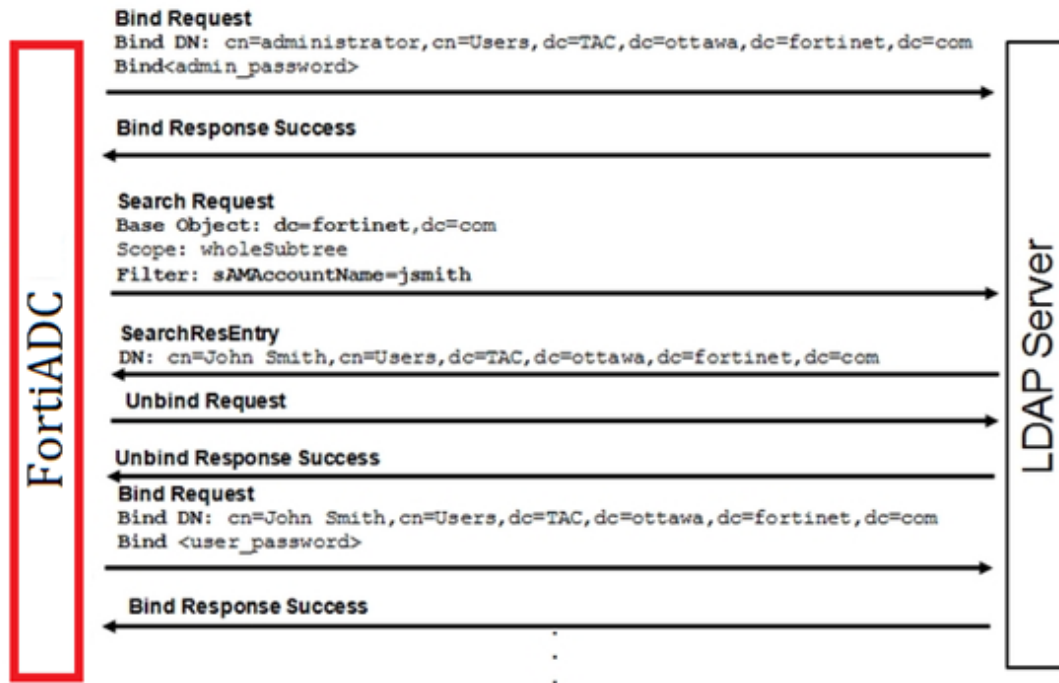
Anonymous bind should be used only if the LDAP server allows it. The LDAP server searches for the client in the entire sub-branches, starting from the specified DN. This bind has two steps: First, FortiADC sends the binding

request to specify the search entry point. Then, it sends a search request with the specified scope and filter to the LDAP server to find the given client.



Regular bind

Regular bind can be used when anonymous binding is not allowed on the LDAP server. Regular bind is similar to anonymous bind. The difference is in the initial step. Unlike anonymous bind, regular bind requires that FortiADC get the access privileges on the LDAP server with the specified user DN in the first step. After it has obtained the authorization, FortiADC can then move on to the second step as it does in anonymous bind.



LDAP over SSL (LDAPS) and StartTLS

LDAP over SSL (LDAPS) and StartTLS are used to encrypt LDAP messages in the authentication process.

LDAPS is a mechanism for establishing an encrypted SSL/TLS connection for LDAP. It requires the use of a separate port, normally Port 636. StartTLS extended operation is an LDAPv3 standard mechanism for enabling TLS (SSL) data confidentiality protection. The mechanism uses an LDAPv3 extended operation to establish an encrypted SSL/TLS connection within an already established LDAP connection.

Configuring LDAP binding

You can use an LDAP authentication server to authenticate administrator or destination server user log-ins.

Basic steps:

1. Configure a connection to an LDAP server that can authenticate administrator or user log-ins.
2. Select the LDAP server configuration when you add administrator users or create user groups.

Before you begin:

- You must know the IP address and port used to access the LDAP server. You must know the CN and DN where user credentials are stored on the LDAP server.
- You must have read-write permission for system settings.

To select an LDAP server:

1. Go to Authentication Management > Remote Server.
2. Select the **LDAP Server** tab.
3. Click **Add** to display the configuration editor.

4. Complete the configuration as described in [Table 80](#).
5. Save the configuration.

Table 80: LDAP server configuration

Settings	Guidelines
Name	<p>Enter the name of LDAP authentication configuration. Valid characters are A-Z, a-z, 0-9, _, and -. No spaces.</p> <p>After you initially save the configuration, you cannot edit the name.</p>
Server	Enter the IP address of the LDAP server.
Port	Enter the port number of the LDAP server. The commonly used port for LDAP is 389.
Common Name Identifier	Enter the Common Name (CN) attribute for the LDAP record. For example: <code>cn</code>
Distinguished Name	<p>If you know the Distinguished Name (DN) of the LDAP server, enter it now. Otherwise, click the Fetch DN button, which enables FortiADC to search for the DN based on the IP address of the LDAP server you've entered above. It automatically populates this field with the DN when it is found.</p> <p>Note: The DN uniquely identifies the device in the LDAP directory. For example: <code>cn=John%20Doe,dc=example,dc=com</code></p>
Fetch DN	Click this button to let FortiADC to fetch the LDAP server's DN. See above.
Bind Type	<ul style="list-style-type: none"> Simple—bind without user search. It can be used only if all the users belong to the same “branch”. Anonymous—bind with user search. It can be used when users are in different “branches” and only if the server allows “anonymous search”. Regular—bind with user search. It can be used when users are in different “branches” and the server does not allow “anonymous search”.
User DN	Available only when Bind Type is “Regular”. In that case, enter the user DN.
Password	Available only when Bind Type is “Regular”. In that case, enter the user password.
Secure Connection	<ul style="list-style-type: none"> Disable LDAPS STARTTLS
CA Profile	This field becomes available only when Secure Connection is set to LDAPS or STARTTLS, regardless of the Bind type being selected. In that case, you can either select a CA that has already been provisioned to secure the connection. You may also leave the field blank if you do not want to secure the connection.

Using a RADIUS authentication server

You can use a RADIUS authentication server to authenticate administrator or destination server user log-ins.

Basic steps:

1. Configure a connection to a RADIUS server that can authenticate administrator or user log-ins.
2. Select the RADIUS server configuration when you add administrator users or user groups.

Before you begin:

- You must know the IP address, port, authentication protocol, and shared secret used to access the RADIUS server.
- You must have read-write permission for system settings.

To create a RADIUS server configuration:

1. Go to Authentication Management > Remote Server.
2. Select the RADIUS Server tab.
3. Click **Add** to display the configuration editor.
4. Complete the configuration as described in [Table 81](#).
5. Save the configuration.

Table 81: RADIUS server configuration

Settings	Guidelines
Name	Configuration name. Valid characters are A-Z, a-z, 0-9, _, and -. No spaces. After you initially save the configuration, you cannot edit the name.
Server	IP address for the server.
Port	Port number for the server. The commonly used port for RADIUS is 1812.
Shared Secret	Shared secret string used when connecting to the server.
Authentication Type	<ul style="list-style-type: none">• PAP—Password authentication protocol.• CHAP—Challenge-Handshake Authentication Protocol.• MS-CHAP—Microsoft version of CHAP.• MS-CHAPv2—Microsoft version of CHAP, version 2.

Using Kerberos Authentication Relay

Kerberos authentication is a computer authentication protocol that works on the basis of tickets (i.e., credentials). It provides several authentication choices, allowing nodes to communicate over a non-secure network to verify each others' identity securely via a Key Distribution Center (KDC) and Service Tickets (STs). It is primarily used for the client-server authentication model and provides mutual authentication by which both the client and the server verify each others' identity.

Kerberos authentication is built upon symmetric key cryptography and requires a trusted third party, and may also resort to the use of public-key cryptography in certain phases of the authentication process. By default, Kerberos Authentication Relay uses UDP port 88.

The Kerberos authentication consists of the following logical components:

- Client
- Authentication Server (AS)
- Ticket Granting Server (TGS)
- Service Server (SS)

Often, the AS and TGS are located on the same physical server, i.e., the KDC.

Authentication Workflow

The following paragraphs highlight the workflow of Kerberos authentication.

Step 1: Client authentication

The client sends a cleartext message of the user ID to the Authentication Server (AS) requesting services on behalf of the user. Note that neither the secret key nor the password is sent to the AS. The AS generates the secret key by hashing the password of the user found at the database, e.g., Active Directory in Windows Server. The AS then checks to see if the client is in its database. If it is, the AS sends back the following two messages to the client:

- Message A: Client/TGS Session Key encrypted using the secret key of the client/user.
- Message B: Ticket Granting Ticket (TGT) which includes the client ID, client network address, ticket validity period, and the client/TGS session key encrypted using the secret key of the TGS.

Once the client receives Messages A and B, it attempts to decrypt Message A with the secret key generated from the password entered by the user. If the user entered password does not match the password in the AS database, the client's secret key will be different and thus unable to decrypt Message A. With a valid password and secret key, the client decrypts Message A to obtain the Client/TGS Session Key. This session key is used for further communications with the TGS. Note that the client cannot decrypt Message B, as it is encrypted using TGS's secret key. At this point, the client has enough information to authenticate itself to the TGS.

Step 2: Client service authorization

When requesting services, the client sends the following messages to the TGS:

- Message C: Composed of the TGT from Message B and the ID of the requested service.
- Message D: Authenticator, which is composed of the client ID and the time-stamp, encrypted using the Client/TGS Session Key.

Upon receiving Messages C and D, the TGS retrieves Message B out of Message C. It decrypts Message B using the TGS secret key. This gives the TGS the "client/TGS session key". Using this key, the TGS decrypts Message D (Authenticator) and sends the following two messages to the client:

- Message E: Client-to-server ticket, which includes the client ID, client network address, validity period, and Client/Server Session Key, encrypted using the service's secret key.
- Message F: Client/Server Session Key encrypted with the Client/TGS Session Key.

Step 3: Client service request

Upon receiving Messages E and F from the TGS, the client has enough information to authenticate itself to the SS. The client connects to the SS and sends the following two messages:

- Message E: From the previous step (the client-to-server ticket, encrypted using the service's secret key).
- Message G: A new Authenticator, which includes the client ID and time-stamp encrypted using the Client/Server Session Key.

The SS decrypts the ticket using its own secret key to retrieve the Client/Server Session Key. Using the sessions key, the SS decrypts the Authenticator and sends the following message to the client to confirm its true identity and willingness to serve the client:

- Message H: The time-stamp found in client's Authenticator, plus 1 in version 4, (but not necessary in version 5[2][3]), encrypted using the Client/Server Session Key.

The client decrypts the confirmation using the Client/Server Session Key and checks whether the time-stamp is correct. If it is correct, then the client can trust the server and start issuing service requests to the server.

The server provides the requested services to the client.

FortiADC Kerberos authentication implementation

Implementation of Kerberos authentication involves the following configurations in FortiADC:

- Authentication Relay. See the following paragraph.
- User Group. See ["Configuring user groups" on page 242](#).
- Authentication Policy. See ["Configuring authentication policies" on page 150](#)
- Virtual Server. See ["Configuring virtual servers" on page 156](#)

Configure Authentication Relay (Kerberos)

Use the following steps to configure Kerberos authentication:

1. Click Authentication Management > Authentication Relay.
2. Click **Add** to open the configuration editor dialog.
3. Make the desired entries or selections as described in the table below.
4. Click **Save** when done.

Table 82: Kerberos authentication configuration

Settings	Guidelines
Name	Specify the name of the authentication relay configuration.
Delegation Type	Select Kerberos.
KDC IP	Enter the IP address of the KDC.
KDC Port	88
Realm	Specify the realm in all upper-case characters.

Settings	Guidelines
Delegator Account	Specify the delegator account. Required.
Delegator Password	Specify the delegator password. Required.
Authorization	Select either of the following: <ul style="list-style-type: none"> HTTP Error 401—If selected, FortiADC relays the authentication credentials only when it encounters an HTTP Error 401 (Unauthorized) from the back-end server. Always—If selected, FortiADC relays the authentication credentials all the time.
Delegated SPN	Specify the delegated SPN. <i>Required.</i>
Domain Prefix Support	Disabled by default. When selected, specify the domain prefix below.
Domain Prefix	Enter the domain prefix only when Domain Prefix Support is enabled. See above.

Using HTTP Basic SSO

When an application uses a Credentials Management API to prompt for user credentials, you must enter the required information that can be validated either by the operating system or by the web application. You can specify your domain credentials information in either of the following formats:

- User Principal Name (UPN)
- Down-Level Logon Name

The UPN format is used to specify an Internet-style name, such as `UserName@Example.Fortinet.com`. The following table presents an anatomy of a UPN:

Table 83: Anatomy of a UPN

Component	Comment	Example
User name	The name of an account	JohnDoell
Separator	The at sign (@)	@
UPN suffix	Also known as the domain name	Example.Fortinet.com

The down-level logon name format specifies a domain and a user account in that domain, for example, `DOMAIN\UserName`. The following table highlights the components of a down-level logon name:

Table 84: Anatomy of a down-level logon name

Component	Description	Example
NetBIOS domain name	Domain name	Domain
Separator	The backslash (\)	\
User account name	Also known as the login name	User name

FortiADC supports HTTP basic SSO when Client Authentication Method is set to be either HTML Form Authentication or HTML Basic Authentication.

For HTTP basic SSO, FortiADC forwards the client's credentials to the web application via the HTTP "Authorization" header. For example, `username/password "user1/fortinet"` from a client is added to the HTTP header in the format `"Authorization: Basic dXNlcjE6Zm9ydGluZXQ="`, and then forwarded to the back-end web application.

You can use either UPN or down-level logon name to log into a web application, and FortiADC adds the domain offload of your logon name for your convenience. Automatically adding the default domain prefix enables you to log in using your user name alone in environments where both user name and domain name are required for the same purpose. This feature comes in handy when you forget your domain name while trying to log into a web application..

Configure HTTP Basic SSO

Use the following steps to configure Kerberos authentication:

1. Click Authentication Management > Authentication Relay.
2. Click **Add** to open the configuration editor dialog.
3. Make the desired entries or selections as described in the table below.
4. Click **Save** when done.

Table 85: Kerberos authentication configuration

Settings	Guidelines
Name	Specify the name of the authentication relay configuration.
Delegation Type	Select HTTP Basic
Authorization	Select either of the following: <ul style="list-style-type: none"> • HTTP Error 401—If selected, FortiADC relays the authentication credentials only when it encounters an HTTP 401 error from the back-end server. • Always—If selected, FortiADC relays the authentication credentials all the time.
Domain Prefix Support	Disabled by default. When selected, you must specify the domain prefix, as described below.

Settings	Guidelines
Domain Prefix	Enter the domain prefix only when you've enabled Domain Prefix Support. See above.

SAML and SSO

Web Single Sign-on (SSO) is an approach that allows single sign-on (SSO) for multiple web applications that have established a common agreement on how to exchange user information. End users provide their credentials only once and are recognized by all of the Web applications, even if they are deployed in different domains and use different identity stores. Web SSO also allows the use of a single identity store by all of the Web apps.

Security Assertion Markup Language (SAML) defines an XML-based framework for describing and exchanging security information among online business entities. It is the most popular protocol for implementing Web SSO.

The SAML protocol has two components—the Service Provider (SP) and the Identify Provider (IDP). They use SAML-defined formatted XML to talk to each other and deliver the identity information called Authentication Assertion.

FortiADC support SAML 2.0, which offers the following benefits:

- Provides support for service provider (SP) and Identity Provider (IDP) Metadata
- Provides single sign-on (SSO) experience for all virtual server resources linked with the user log-in

Functioning as an SP, FortiADC supports the following IDPs:

- FortiAuthenticator (Factory default)
- Shibboleth
- OpenAM/OpenSSO

Configure a SAML service provider

You must configure your SPs in order to use SAML authentication. To configure an SP, you must have the required IDP metadata file imported into FortiADC ahead of time. See ["Import IDP Metadata" on page 256](#) for more information.

Once you have imported the needed IDP metadata file into FortiADC, you can use the following steps to configure a SAML service provider:

1. Click Authentication Management > SAML.
2. Select the SAML Service Providers tab, if it is not selected.
3. Click Add to open the SAML Service Providers configuration editor.
4. Make the desired entries or selections, as described in the table below.
5. Click Save when done.

Table 86: Configure a SAML service provider

Parameter	Description
SAML Service Providers	Use this page to configure an SAML service provider.
Name	Specify a unique name for the SAML service provider.
Entity ID	Specify the SAML service provider's entity ID, which is the SAML service provider's URL.
Local Certification	Select an option. The default is Factory.
Service URL	/SSO
Assertion Con- suming Service Bind- ing Type	Post.
Assertion Con- suming Service Path	/SAML2/Post
Single Logout Bind- ing Type	Post
Single Logout Path	/SLO/Logout
IDP Metadata	Select an IDP metadata file. Note: You must have the IDP metadata file imported into FortiADC ahead of time.
Metadata Export Ser- vice Location	/Metadata
Authentication Ses- sion Lifetime	28800
Authentication Ses- sion Timeout	3600
SSO Status	Enable(d) by default, which allows FortiADC to forward SSO information to the real server, which in turn gets the authentication information and implements the SSO function.
Export Assertion Status	Enabled by default, which allows FortiADC to send authentication assertions (i.e., identity information) to the real server that requests the information.
Export Assertion Path	/GetAssertion

Parameter	Description
Export Cookie Status	Enable(d) by default, which allows FortiADC to send to the real server the cookie of a site that the user last visited.
Export Assertion ACL	
IP Netmask	Enter the IP address of the real server (or the IP Netmask if the real server is one of a group of real servers) that requests authentication assertions.

Import IDP Metadata

A SAML metadata file provides the information of a client, such as its entity ID, credential, and so on. It also contains a of couple of URLs so that the server knows where to send different requests, e.g., log-in requests, attribute query requests, etc. You need to import this metadata to your SAML component so that it knows which client it should talk to.

Another purpose is to establish a trust relationship between the Service Provider (SP) and Identity Provider (IdP). In this case, SAML metadata is used to exchange configuration information between the SP and the IdP, and vice versa. The metadata can be signed and encrypted so that the data is transferred securely. The other side may need the corresponding public key to validate and decrypt it and then can be used to understand and establish the connection with the SP or IdP

To import a SAML IDP metadata file:

1. Click Authentication Management > SAML.
2. Select the IDP Metadata tab.
3. Click Import.
4. Follow the instructions onscreen to import the IDP metadata file.

Chapter 10: Shared Resources

This chapter includes the following topics:

- [Configuring health checks](#)
- [Creating schedule groups](#)
- [Creating IPv4 address objects](#)
- [Configuring IPv4 address groups](#)
- [Configuring IPv6 address objects](#)
- [Configuring IPv6 address groups](#)
- [Managing the ISP address books](#)
- [Creating service objects](#)
- [Creating service groups](#)

Configuring health checks

In server load balancing deployments, the system uses health checks to poll the members of the real server pool to test whether an application is available. You can also configure additional health checks to poll related servers, and you can include results for both in the health check rule. For example, you can configure an HTTP health check test and a RADIUS health check test. In a web application that requires user authentication, the web server is deemed available only if the web server and the related RADIUS server pass the health check.

In link load balancing deployments, the health check can poll either the ISP link group member itself or a “beacon” server that is deployed on the other side of the ISP link. A beacon is an IP address that must be reachable in order for the link to be deemed available. A beacon can be any IP address, such as a main office, core router, or virtual server at another data center.



If you expect a backend server is going to be unavailable for a long period, such as when it is undergoing hardware repair, it is experiencing extended down time, or when you have removed it from the server farm, you can improve the performance of the FortiADC system by setting the status of the pool member to Disabled, rather than allowing the system to continue to attempt health checks.

Table 87 describes the predefined health checks. You can get started with these or create custom objects.

Table 87: Predefined health check configuration objects

Predefined	Description
LB_HLTHCK_HTTP	Sends a HEAD request to the server port 80. Expects the server to return an HTTP 200.
LB_HLTHCK_HTTPS	Sends a HEAD request to the server port 443. Expects the server to return an HTTP 200.

Predefined	Description
LB_HLTHCK_ICMP	Pings the server.
LB_HLTHCK_TCP_ECHO	Sends a TCP echo to server port 7. Expects the server to respond with the corresponding TCP echo.

Before you begin:

- You must have a good understanding of TCP/IP and knowledge of the services running on your backend servers.
- You must know the IP address, port, and configuration details for the applications running on backend servers. For some application protocol checks, you must specify user credentials.
- You must have Read-Write permission for Load Balance settings.

After you have configured a health check, you can select it in the SLB server pool, LLB link group, or GLB server configuration.

To configure a health check:

1. Go to Shared Resources > Health Check.
2. Click **Add** to display the configuration editor.
3. Select one of the following options:
 - ICMP
 - TCP Echo
 - TCP
 - HTTP
 - HTTPS
 - DNS
 - RADIUS
 - SMTP
 - POP3
 - IMAP4
 - RADIUS Accounting
 - FTP
 - TCP Half Open Connection
 - TCP SSL
 - SNMP
 - SSH
 - L2 Detection
 - UDP
 - SIP
 - SIP-TCP
 - SNMP-Custom
 - RSTP
 - MySQL
4. Complete the configuration as described in [Table 88](#).
5. Save the configuration.



You can clone a predefined configuration object to help you get started with a user-defined configuration.


To clone a configuration object, click the clone icon  that appears in the tools column on the configuration summary page.

Table 88: Health check configuration

Settings	Guidelines
General	
Name	Configuration name. Valid characters are A-Z, a-z, 0-9, _, and -. No spaces. After you initially save the configuration, you cannot edit the name.
Type	Select a type of health check.
Destination Address Type	<ul style="list-style-type: none"> • IPv4 • IPv6
Destination Address	<p>IP address to send health check traffic.</p> <p>In server load balancing deployments, if you do not specify an IP address, the real server IP address is used. You might configure IP address for a health check if you are configuring a combination of health checks to poll related servers.</p> <p>In link load balancing deployments, if you do not specify an IP address, the destination IP address is the address of the gateway. You can configure IP address if you want to test connectivity to a beacon on the other side of the gateway, or if you want to test whether service traffic is allowed to pass through the link.</p>
Hostname	For HTTP or HTTPS health checks, you can specify the hostname (FQDN) instead of the destination IP address. This is useful in VM environments where multiple applications have the same IP address.
Interval	Seconds between each health check. Should be more than the timeout to prevent overlapping health checks. The default is 10.
Timeout	Seconds to wait for a reply before assuming that the health check has failed. The default is 5.
Up Retry	Attempts to retry the health check to confirm server availability. The default is 1.
Down Retry	Attempts to retry the health check to see if a down server has become available. The default is 1.
Specifics	
ICMP	
No specific options	Simple ping to test connectivity.
TCP Echo	

Settings	Guidelines
No specific options	Simple ping to test connectivity.
TCP / TCP Half Open Connection / UDP	
Port	Listening port number of the backend server. Usually HTTP is 80, FTP is 21, DNS is 53, POP3 is 110, IMAP4 is 143, RADIUS is 1812, and SNMP is 161.
TCP SSL	
Port	Listening port number of the backend server. Usually HTTP is 80, FTP is 21, DNS is 53, POP3 is 110, IMAP4 is 143, RADIUS is 1812, and SNMP is 161.
SSL Ciphers	Default selections are recommended.
Local Cert	For TCP SSL only. Click the down arrow and select a local SSL Health Check Client certificate from the list menu. The certificate titled "Factory" is the default certificate shipped with your FortiADC. The rest, if any, are the custom certificates that you have created.
HTTP/HTTPS	
Port	Listening port number of the backend server. Usually HTTP is 80. If testing an HTTP proxy server, specify the proxy port.
SSL Ciphers	For HTTPS only. Default selections are recommended.
Local Cert	For HTTPS only. See TCP / TCP Half Open Connection / TCP SSL / UDP above.
HTTP CONNECT	<p>If the real server pool members are HTTP proxy servers, specify an HTTP CONNECT option:</p> <ul style="list-style-type: none"> • Local CONNECT—Use HTTP CONNECT to test the tunnel connection through the proxy to the remote server. The member is deemed available if the request returns status code 200 (OK). • Remote CONNECT—Use HTTP CONNECT to test both the proxy server response and remote server application availability. If you select this option, you can configure an HTTP request within the tunnel. For example, you can configure an HTTP GET/HEAD request to the specified URL and the expected response. • No CONNECT—Do not use the HTTP CONNECT method. This option is the default. The HTTP CONNECT option is useful to test the availability of proxy servers only. <p>See the FortiADC Deployment Guide for FortiCache for an example that uses this health check.</p>

Settings	Guidelines
Remote Host	If you use HTTP CONNECT to test proxy servers, specify the remote server IP address.
Remote Port	If you use HTTP CONNECT to test proxy servers, specify the remote server port.
Method Type	HTTP method for the test traffic: <ul style="list-style-type: none"> • HTTP GET—Send an HTTP GET request to the server. A response to an HTTP GET request includes HTTP headers and HTTP body. • HTTP HEAD—Send an HTTP HEAD request. A response to an HTTP HEAD request includes HTTP headers only.
Send String	The request URL, such as /contact.php.
Receive String	A string expected in return when the HTTP GET request is successful.
Status Code	The health check sends an HTTP request to the server. Specify the HTTP status code in the server reply that indicates a successful test. Typically, you use status code 200 (OK). Other status codes indicate errors.
Match Type	What determines a failed health check? <ul style="list-style-type: none"> • Match String • Match Status • Match All (match both string and status) Not applicable when using HTTP HEAD. HTTP HEAD requests test status code only.
DNS	
Domain Name	The FQDN, such as www.example.com, to use in the DNS A/AAAA record health check.
Address Type	<ul style="list-style-type: none"> • IPv4 • IPv6
Host Address	IP address that matches the FQDN, indicating a successful health check.
RADIUS / RADIUS Accounting	
Port	Listening port number of the backend server. Usually RADIUS is 1812 and RADIUS accounting is 1813.
Username	User name of an account on the backend server.
Password	The corresponding password.

Settings	Guidelines
Password Type	<ul style="list-style-type: none"> • User—If the backend server does not use CHAP, select this option. • CHAP—If the backend server uses CHAP and does not require a secret key, select this option.
Secret Key	The secret set on the backend server.
NAS IP Address	NAS IP address RADIUS attribute (if the RADIUS server requires this attribute to make a connection).
SIP / SIP-TCP	
Port	Specify the port number. Valid values range from 0 to 65535.
SIP Request Type	Specify the SIP request type to be used for health checks: <ul style="list-style-type: none"> • SIP Options • SIP Register
Status Code	The expected response code. If not set, response code 200 is expected. Specify 0 if any reply should indicate the server is available.
SMTP	
Port	Listening port number of the backend server. Usually SMTP is 25.
Domain Name	The FQDN, such as www.example.com, to use in the SMTP HELO request used for health checks. If the response is OK (250), the server is considered as up. If there is error response (501) or no response at all, the server is considered down.
POP3	
Port	Listening port number of the backend server. Usually POP3 is 110.
Username	User name of an account on the backend server.
Password	The corresponding password.
IMAP4	
Port	Listening port number of the backend server. Usually IMAP4 is 143.
Username	User name of an account on the backend server.
Password	The corresponding password.

Settings	Guidelines
Folder	Select an email mailbox to use in the health check. If the mailbox does not exist or is not accessible, the health check fails. The default is INBOX.
FTP	
Port	Listening port number of the backend server. Usually FTP is 21.
User name	User name of an account on the backend server.
Password	The corresponding password.
File	Specify a file that exists on the backend server. Path is relative to the initial login path. If the file does not exist or is not accessible, the health check fails.
Passive	Select this option if the backend server uses passive FTP.
SNMP	
Port	Listening port number of the backend server. Usually SNMP is 161 or 162.
CPU	Maximum normal CPU usage. If overburdened, the health check fails.
Memory	Maximum normal RAM usage. If overburdened, the health check fails.
Disk	Maximum normal disk usage. If the disk is too full, the health check fails.
Agent type	<ul style="list-style-type: none"> • UCD • Windows 2000
Community	Must match the SNMP community string set on the backend server. If this does not match, all SNMP health checks fail.
Version	SNMP v1 or v2c.
CPU Weight	100
Memory Weight	100
Disk Weight	100
SNMP-Custom	
Port	Listening port number of the backend server. Usually SNMP is 161 or 162.

Settings	Guidelines
Community	Must match the SNMP community string set on the backend server. If this does not match, all SNMP health checks fail.
Version	SNMP v1 or v2c.
OID	String specifying the OID to query
Value Type	Abstract syntax notation (ASN) value type: <ul style="list-style-type: none"> • ASN_INTEGER • ASN_OCTET_STR • ASN_OBJECT_ID • ASN_COUNTER • ASN_UINTeger
Compare Type	<ul style="list-style-type: none"> • Equal • Less • Greater
Counter Value	Specify the value for the evaluation.
SSH	
Port	Listening port number of the backend server. Usually SSH is 22.
Username	Username for test login.
Password	Corresponding password.
L2 Detection	
No specific options	Link Layer health checker. Sends ARP (IPv4) or NDP (IPv6) packets to test whether a physically connected system is available.
RTSP	
Port	Specify the listening port number. Valid values range from 0 to 65535.
RTSP Method Type	RTSP Options
Status Code	200
MySQL	

Settings	Guidelines
Port	Specify the listening port number of the MySQL server. Valid values range from 0 to 65535.
Username	Specify the database user name. (Optional)
Password	Specify the database password, if applicable.
MySQL Server Type	Select either of the following: <ul style="list-style-type: none"> • Master (Default) • Slave



In SLB deployments, a health check port configuration specifying port 0 acts as a wildcard. The port for health check traffic is imputed from the real server pool member.

In LLB and GLB deployments, specifying port 0 is invalid because there is no associated configuration to impute a proper port. If your health check port configuration specifies port 0, you will not be able to use it in an LLB or GLB configuration.

Monitoring health check status

FortiADC enables you to monitor the health of server in real time directly from your desktop, as described below.

1. Click Shared Resources > Health Check.
2. Click the **Health Check Monitor** tab.
3. Configure the health check monitor as described in the table below.
4. Click **Start** to perform the health check. The result will show in the Monitor Information.

Table 89: Checking server health

Parameter	Description
IP Address	Enter the IP address of the remote server.
Health Check	Select the health check configuration.
Port	Enter the port number, if applicable. Note: This field is available only for health check configurations that require port numbers.

Creating schedule groups

You create schedule objects to use in link load balancing policies. A policy rule can be time-bound: one time, daily, weekly, or monthly.

Basic Steps

1. Create a schedule object.
2. Select the schedule when you configure the link policy.

Before you begin:

- You must have Read-Write permission for System settings.

To create schedule objects:

1. Go to Shared Resources > Schedule Group.
2. Click **Add** to display the configuration editor.
3. Give the schedule a name, save it, and add schedule members as described in [Table 90](#).
4. Save the configuration.

Table 90: Schedule member configuration

Settings	Guidelines
Name	Unique group name. Valid characters are A-Z, a-z, 0-9, _, and -. No spaces. After you initially save the configuration, you cannot edit the name.
Member	
Name	Unique member name. Valid characters are A-Z, a-z, 0-9, _, and -. No spaces. After you initially save the configuration, you cannot edit the name.
Type	<ul style="list-style-type: none"> • One Time • Daily • Weekly • Monthly
Start Date	YYYY/MM/DD.
End Date	YYYY/MM/DD.
Start Time	HH:MM.
End Time	HH:MM.

Creating IPv4 address objects

You create address objects to specify matching source and destination addresses in policies.

The following policies use address objects:

- Firewall policies
- QoS policies
- Connection limit policies
- Link load balancing policies

Note: For link load balancing, you can also add address objects to address groups, which can then be used in link load balance policies.

Basic Steps

1. Create address objects.
2. Select them when you configure address groups or policies.

Note: Before you begin, you must have Read-Write permission for System settings.

To create an address object:

1. Click Shared Resources > Address.
2. Click **Add** to display the configuration editor.
3. Complete the configuration as described in [Table 91](#).
4. Click **Save**.

Table 91: Address object configuration

Settings	Guidelines
Name	Configuration name. Valid characters are A-Z, a-z, 0-9, _, and -. No spaces. After you initially save the configuration, you cannot edit the name.
Type	<ul style="list-style-type: none"> • IPv4/Netmask • Address Range
IPv4/Netmask (or IPv6/Netmask)	Specify a subnet using the IP address/mask notation.
Address Range	Specify the start and end of an address range.

Configuring IPv4 address groups

You configure address group objects when you have more than one address object you want to specify in rules that match source or destination addresses. For example, if you subscribe customer 1 and customer 2 to a group of links, then you can create rules that match the customer 1 OR customer 2 address space and load balance the set of gateways assigned to them.

The following policies use address groups:

- Link load balancing policies

Basic Steps

1. Create address objects.
2. Configure address group objects. You can add up to 256 members in a group.
3. Select the address groups when you configure your policies.

Before you begin:

- You must have Read-Write permission for System settings.

To configure an address group:

1. Click Shared Resources > Address.
2. Click the **Address Group** tab.
3. Click **Add** to display the configuration editor.
4. Complete the configuration as described in [Table 92](#).
5. Click **Save**.

Table 92: Address Group configuration

Settings	Guidelines
Name	Configuration name. Valid characters are A-Z, a-z, 0-9, _, and -. No spaces. After you initially save the configuration, you cannot edit the name.
Member List	
Name	Configuration name. Valid characters are A-Z, a-z, 0-9, _, and -. No spaces. After you initially save the configuration, you cannot edit the name.
Address	Select an address object.

Creating IPv6 address objects

You create address objects to specify matching source and destination addresses in policies.

The following policies use address objects:

- Firewall policies
- QoS policies
- Connection limit policies
- Link load balancing policies

Note: For link load balancing, you can also add address objects to address groups, which can then be used in link load balance policies.

Basic Steps

1. Create address objects.
2. Select them when you configure address groups or policies.

Note: Before you begin, you must have Read-Write permission for System settings.

To create an address object:

1. Click Shared Resources > IPv6 Address.
2. Click **Add** to display the configuration editor.
3. Complete the configuration as described in [Table 93](#).
4. Click **Save**.

Table 93: IPv6 Address object configuration

Settings	Guidelines
Name	Configuration name. Valid characters are A-Z, a-z, 0-9, __, and -. No spaces. After you initially save the configuration, you cannot edit the name.
Type	<ul style="list-style-type: none"> • IPv6/Netmask • Address Range
IPv4/Netmask (or IPv6/Netmask)	Specify a subnet using the IP address/mask notation.
Address Range	Specify the start and end of an address range.

Configuring IPv6 address groups

You configure address group objects when you have more than one address object you want to specify in rules that match source or destination addresses. For example, if you subscribe customer 1 and customer 2 to a group of links, then you can create rules that match the customer 1 OR customer 2 address space and load balance the set of gateways assigned to them.

The following policies use address groups:

- Link load balancing policies

Basic Steps

1. Create address objects.
2. Configure address group objects. You can add up to 256 members in a group.
3. Select the address groups when you configure your policies.

Before you begin:

- You must have Read-Write permission for System settings.

To configure an address group:

1. Click Shared Resources > Address.
2. Click the **IPv6 Address Group** tab.
3. Click **Add** to display the configuration editor.
4. Complete the configuration as described in [Table 94](#).
5. Click **Save**.

Table 94: Address Group configuration

Settings	Guidelines
Name	Configuration name. Valid characters are A-Z, a-z, 0-9, _, and -. No spaces. After you initially save the configuration, you cannot edit the name.
Member List	
Name	Configuration name. Valid characters are A-Z, a-z, 0-9, _, and -. No spaces. After you initially save the configuration, you cannot edit the name.
Address	Select an address object.

Managing ISP address books

ISP address books contain IP subnet addresses and associated province location settings for ISP links.

The following policies use the ISP address book objects:

- ISP routes
- LLB proximity routes
- LLB policies
- GLB data center configuration

The province setting is used in GLB deployments in China to enable location awareness that is province-specific. For example, a user can be directed to a data center in specific location inside the country, such as Beijing or Guangdong, rather than simply China.

[Figure 49](#) shows the three types of address book entries:

- **Predefined**—Addresses and associated province location settings for China Mobile, China Telecom, and China Unicom. The IP subnet addresses in the predefined address books are not exposed in the user interface. The predefined package is provided to make it easier for you to configure a route when all you know and all you need to know is the name of the ISP that hosts the link.
- **Restored**—Addresses imported from a text file. The IP subnet addresses in the restored address books are not exposed in the user interface. “Restored” addresses can help you rapidly build an ISP address book configuration. “Restored” addresses can help you rapidly build an ISP address book configuration.
- **User-defined**—In the ISP address configuration, you can modify the predefined and restored address books by specifying subnets to add or exclude from them. This gives you flexibility in case you encounter address conflicts or the ISP instructs you to add a subnet address manually.

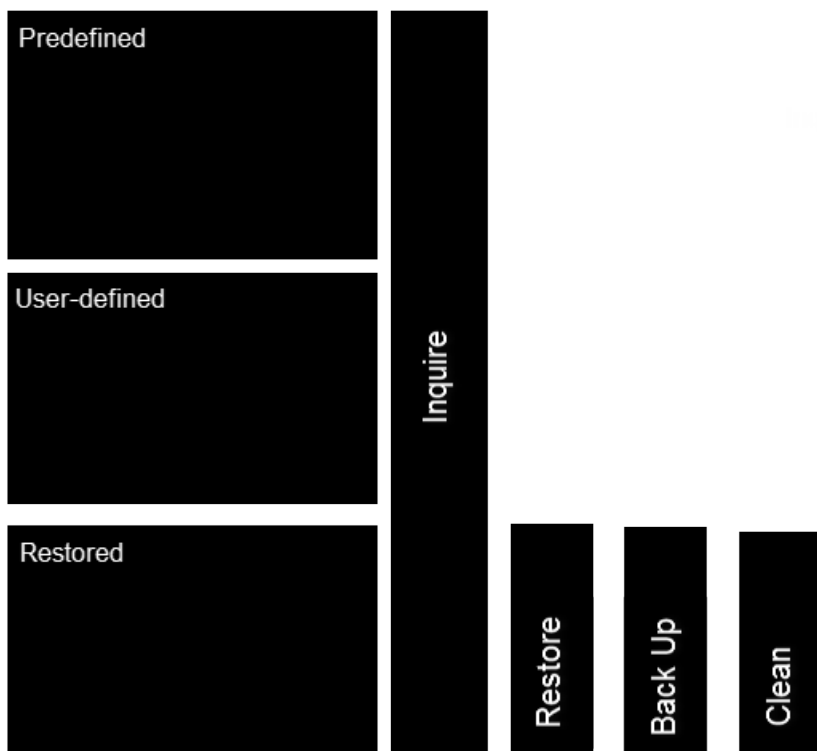
You can also create new user-defined entries for other ISPs.

Note: In systems with multiple VDOMs, these commands apply to the current VDOM only. In other words, if you configure an exclusion, it is applicable to the current VDOM only; it does not change the predefined address book.

You can use the **Inquire** utility to see whether an IP address belongs to any of the address books. If an address can be found in more than one address book, the results are returned in the following priority:

1. User-defined
2. Restored
3. Predefined

Figure 49: ISP address book types



The text file for the Restored entries has the following format:

```
#this is a comment line
ISP name:ABC
Province:Beijing
```

```

1.1.1.0/24
Province:Unknown
2.2.0.0 255.255.0.0
#this is a comment line too
3.3.3.3/32
ISP name:DEF
Province:Shanghai
4.4.4.0 255.255.255.0
5.5.0.0/16

```

You use the **Restore** utility to import the file and the **Back Up** utility to export it.

You use the **Clean** utility to erase entries that were imported from the text file. The clean operation does not affect the predefined addresses or user-configured entries. If a restored entry has user-configured elements (for example, an exclude list), the clean operation clears the addresses but preserves the configuration and converts it to a user-defined type.

Basic Steps

1. Create ISP address objects.
2. Select them when you configure your policies.

Note: Before you begin, you must have read-write permission for System settings.

Create an ISP address book object

To create an ISP address book object:

1. Click Shared Resource > Address.
2. Click the **ISP Address** tab.
3. Click **Add**. The ISP Address dialog opens.
4. Complete the configuration as described in [Table 95](#).
5. Click **Save**.

Table 95: ISP address object configuration

Settings	Guidelines
Name	Configuration name. Valid characters are A-Z, a-z, 0-9, _, and -. No spaces. After you initially save the configuration, you cannot edit the name.
Address	Address/mask notation specifying a subnet to add it to the address book entry.
Excluded Address	Address/mask notation specifying a subnet to be excluded from the address book entry. Create exclusions to predefined and restored address books only. Note: This field applies to predefined and restored address books only; it is not applicable or available for user-defined address books.

Settings	Guidelines		
Province	Select the associated province location. The configuration supports the following selections:		
	Anhui	Henan	Shanxi (Taiyuan)
	Beijing	Hubei	Shanxi (Xian)
	Chongqing	Hunan	Sichuan
	Fujian	Jiangsu	Tianjin
	Gansu	Jiangxi	Xianggang
	Guangdong	Jilin Liaoning	Xinjiang
	Guangxi	Neimenggu	Xizang
	Guizhou	Ningxia	Yunnan
	Hainan	Qinghai	Zhejiang
	Hebei	Shandong	Unknown
	Heilongjiang	Shanghai	

Creating service objects

FortiADC provides more than two dozen predefined services, as shown on the Shared Resources > Service > Service page. In addition, it allows you to create your service objects as well. Service objects are an important part of the following policy configurations:

- Firewall policies
- QoS policies
- Connection limit policies
- Link load balancing policies

Note: For link load-balancing, you can also add service objects to service groups; then use service groups in LLB policies.

Basic Steps

1. Create service objects.
2. Select them when you configure service groups or policies.

Before you begin:

- You must have Read-Write permission for System settings.

To create a service object:

1. Go to Shared Resources > Service.
2. Select the **Service** tab.

3. Click **Add** to display the configuration editor.
4. Complete the configuration as described in [Table 96](#).
5. Save the configuration.

Table 96: Service object configuration

Settings	Guidelines
Name	Configuration name. Valid characters are A-Z, a-z, 0-9, _, and -. No spaces. Note: Once created, the name cannot be changed.
Protocol Type	Select one of the following: <ul style="list-style-type: none"> • ip (default) • icmp • tcp • udp • tcp-and-udp • sctp
Protocol	1 Note: This applies only when Protocol Type is set to IP. In that case, it displays the protocol number without port.
Specify Source Port	This option becomes available when TCP, UDP, SCTP, or TCP-AND-UDP is selected as the protocol type. When selected, you also need to specify the Minimum Source Port and Maximum Source Port below.
Minimum Source Port	1
Maximum Source Port	65535
Minimum Destination Port	1
Maximum Destination Port	-65535

Creating service groups

You configure service group objects when you have more than one service you want to specify in a rule that matches service. You can group all Web services and group all mail services, for example, if you want to have

rules that treat those as groups.

The following policies use service groups:

- Link load balancing policies

Basic Steps

1. Create service objects.
2. Configure service group objects. You can add up to 256 members in a group.
3. Select the service groups when you configure your policies.

Before you begin:

- You must have Read-Write permission for System settings.

To configure a service group:

1. Go to Shared Resources > Service.
2. Click **Service Group**.
3. Click **Add** to display the configuration editor.
4. Complete the configuration as described in [Table 97](#).
5. Save the configuration.

Table 97: Service Group configuration

Settings	Guidelines
Name	Specify a unique name for the service group configuration. Valid characters are A-Z, a-z, 0-9, _, and -. No spaces. After you initially save the configuration, you cannot edit the name.
Member List	
Name	Configuration name. Valid characters are A-Z, a-z, 0-9, _, and -. No spaces. After you initially save the configuration, you cannot edit the name.
Service	Select a service object.

Chapter 11: Basic Networking

This chapter includes the following topics:

- [Configuring network interfaces](#)
- [Configuring static routes](#)
- [Configuring policy routes](#)

See [Chapter 17: Advanced Networking](#) for advanced topics.

Configuring network interfaces

This topic includes the following information:

- [Using physical interfaces](#)
- [Using VLAN interfaces](#)
- [Using aggregate interfaces](#)
- [Configuring network interfaces](#)

Using physical interfaces


Each physical network port (or, on FortiADC-VM, a vNIC) has a network interface that directly corresponds to it—that is, a “physical network interface.”

Physical ports have three uses:

- Management—The network interface named port1 is typically used as the management interface.
- HA—If you plan to deploy HA, you must reserve a physical port for HA heartbeat and synchronization traffic. Do *not* configure the network interface that will be used for HA; instead, leave it unconfigured or “reserved” for HA.
- Traffic—The remaining physical ports can be used for your target traffic—these are your “traffic interfaces.”

Traffic interfaces can be associated with logical interfaces. The system supports two types of logical interfaces: VLAN and aggregate. [Figure 50](#) illustrates how physical ports are associated with physical and logic interfaces.

Figure 50: Physical and logical interfaces



With VLANs, multiple VLAN logical interfaces are associated with a single physical port. With link aggregation, it is the reverse: multiple physical interfaces are associated with a single aggregate logical interface.

Table 98 lists factory default IP addresses for physical network interfaces.

Table 98: Physical network interfaces

Network Interface*	IPv4 Address/Netmask	IPv6 Address/Netmask
port1	192.168.1.99/24	::/0
port2	0.0.0.0/0	::/0
port3	0.0.0.0/0	::/0
port4	0.0.0.0/0	::/0
...		
* The number of physical network interfaces varies by model.		

Using VLAN interfaces

You can use [IEEE 802.1q](#) VLAN to reduce the size of a broadcast domain, thereby reducing the amount of broadcast traffic received by network hosts, improving network performance.

Unlike physical LANs, VLANs do not require you to install separate hardware switches and routers to achieve this effect. Instead, VLAN-compliant switches restrict broadcast traffic based upon whether its VLAN ID matches that of the destination network. As such, VLAN trunks can be used to join physically distant broadcast domains as if they were close.

The VLAN ID is part of the tag that is inserted into each Ethernet frame in order to identify traffic for a specific VLAN. FortiADC appliances handle VLAN header addition automatically, so you do not need to adjust the maximum transmission unit (MTU). Depending on whether the device receiving a packet operates at Layer 2 or Layer 3 of the network, a VLAN tag might be added, removed, or rewritten before forwarding to other nodes on the network. For example, a Layer 2 switch typically adds or removes a tag when forwarding traffic among members of the VLAN, but does not route tagged traffic to a different VLAN ID. In contrast, a FortiADC content-based routing policy might forward traffic between different VLAN IDs (also known as inter-VLAN routing).

Cisco Discovery Protocol (CDP) is supported for VLANs.

Note: VLANs are not designed to be a security measure, and should not be used where untrusted devices and/or individuals outside of your organization have access to the equipment. VLAN tags are not authenticated, and can be ignored or modified by attackers. VLAN tags rely on the voluntary compliance of the receiving host or switch.

Using aggregate interfaces

Link aggregation (also called NIC teaming/bonding or link bundling) forms a network interface that queues and transmits over multiple wires (also called a port channel), instead of only a single wire (as FortiADC would normally do with a single network interface per physical port). This multiplies the bandwidth that is available to the network interface, and therefore is useful if FortiADC is deployed inline with your network backbone.

Link aggregation on FortiADC complies with [IEEE 802.1ax](#) and [IEEE 802.3ad](#) and distributes Ethernet frames using a modified round-robin behavior. If a port in the aggregation fails, traffic is redistributed automatically to the remaining ports with the only noticeable effect being a reduced bandwidth. When broadcast or multicast traffic is received on a port in the aggregation, reverse traffic will return on the same port.

When link aggregation uses a round-robin that considers only Layer 2, Ethernet frames that belong to an HTTP request can sometimes arrive out of order. Because network protocols at higher layers often do not gracefully handle this (especially TCP, which may decrease network performance by requesting retransmission when the expected segment does not arrive), FortiADC's frame distribution algorithm is configurable. For example, if you notice that performance with link aggregation is not as high as you expect, you could try configuring FortiADC to queue related frames consistently to the same port by considering the IP session (Layer 3) and TCP connection (Layer 4), not simply the MAC address (Layer 2).

You must also configure the router, switch, or other link aggregation control protocol (LACP)-compatible device to which FortiADC is connected with the same speed/duplex settings, and it must have ports that can be aggregated. In a deployment like this, the two devices use the cables between the ports to form a trunk, not an accidental Layer 2 (link) network loop. FortiADC uses LACP to detect the following conditions:

- Suitable links between itself and the other device, and form a single logical link.
- Individual port failure so that the aggregate can redistribute queuing to avoid a failed port.

Configuring network interfaces

You can edit the physical interface configuration. You cannot create or delete a physical interface configuration.

Before you begin:

- You must have Read-Write permission for System settings.

To configure a network interface:

1. Go to Networking > Interface.
2. Double-click the row for a physical interface to edit its configuration or click **Add** if you want to configure an aggregate or VLAN interface.
3. Complete the configuration as described in [Table 99](#).
4. Save the configuration.

Table 99: Network interface configuration

Settings	Guidelines
Common Settings	
Name	Configuration name. Valid characters are A-Z, a-z, 0-9, _, and -. No spaces. After you initially save the configuration, you cannot edit the name
Status	The Status column is not the detected physical link status; it is the administrative status (Up/Down) that indicates whether you permit the network interface to receive and/or transmit packets.

Settings	Guidelines
Allow Access	<p>Allow inbound service traffic. Select from the following options:</p> <ul style="list-style-type: none"> • HTTP—Enables connections to the web UI. We recommend this option only for network interfaces connected to a trusted private network, or directly to your management computer. • HTTPS—Enables secure connections to the web UI. We recommend this option instead of HTTP. • Ping—Enables ping and traceroute to be received on this network interface. When it receives an ECHO_REQUEST (“ping”), FortiADC will reply with ICMP type 0 (ECHO_RESPONSE or “pong”). • SNMP—Enables SNMP queries to this network interface. • SSH—Enables SSH connections to the CLI. We recommend this option instead of Telnet. • Telnet—Enables Telnet connections to the CLI. We recommend this option only for network interfaces connected to a trusted private network, or directly to your management computer.
Dedicated HA management IP	<p>If selected, the interface will be reserved as an HA management interface. Once this management interface is reserved, you can configure a different IP address, administrative access and other settings for this interface for each cluster unit. Then by connecting this interface of each cluster unit to your network, you can manage each cluster unit separately from a different IP address.</p>
Virtual Domain	<p>If applicable, select the virtual domain to which the configuration applies.</p>
Mode	<ul style="list-style-type: none"> • Static—Specify a static IP address. The IP address must be on the same subnet as the network to which the interface connects. Two network interfaces cannot have IP addresses on the same subnet (i.e. overlapping subnets). • PPPoE—Use PPPoE to retrieve a configuration for the IP address, gateway, and DNS server. For example, if this interface uses a DSL connection to the Internet, your ISP may require this option.
Static	
Traffic Group	<p>Select either of the following:</p> <ul style="list-style-type: none"> • Default • Create New
Floating	<p>Enable/Disable floating IP.</p>
Floating IP	<p>Enter the floating IP.</p>
IPv4/Netmask	<p>Specify the IP address and CIDR-formatted subnet mask, separated by a forward slash (/), such as 192.0.2.5/24. Dotted quad formatted subnet masks are not accepted.</p>

Settings	Guidelines
IPv6/Netmask	Specify the IP address and CIDR-formatted subnet mask, separated by a forward slash (/), such as 2001:0db8:85a3::8a2e:0370:7334/64. Dotted quad formatted subnet masks are not accepted.
Secondary IP Address	<p>Secondary IP addresses can be used when you deploy the system so that it belongs to multiple logical subnets. If you assign multiple IP addresses to an interface, you must assign them static addresses.</p> <p>To add secondary IP addresses, enable the feature and save the configuration. After you have saved it the first time, you can edit it to add secondary IP addresses and enable inbound traffic to that address.</p>
PPPoE	
Username	PPPoE account user name.
Password	PPPoE account password.
Discovery Retry Timeout	Seconds the system waits before it retries to discover the PPPoE server. The default is 5 seconds. The valid range is 1-255.
DNS Server Override	Use the DNS addresses retrieved from the PPPoE server instead of the one configured in the FortiADC system settings.
Retrieve Default Gateway	Use the default gateway retrieved from the PPPoE server instead of the one configured in the FortiADC system settings.
Type	<p>If you are editing the configuration for a physical interface, you cannot set the type. If you are configuring a logical interface, you can select from the following options:</p> <ul style="list-style-type: none"> Aggregate—A logical interface you create to support the aggregation of multiple physical interfaces. VLAN—A logical interface you create to VLAN subinterfaces on a single physical interface.
Aggregate	
Member	Select the physical interfaces that are included in the aggregation.
Aggregate Mode	<p>Link aggregation type:</p> <ul style="list-style-type: none"> 802.3ad Balance-alb Balance-rr Balance-tlb Balance-xor Broadcast

Settings	Guidelines
Aggregate Algorithm	<p>Connectivity layers that will be considered when distributing frames among the aggregated physical ports:</p> <ul style="list-style-type: none"> • Layer 2 • Layer 2-3 • Layer 3-4
VLAN	
VLAN ID	<p>VLAN ID of packets that belong to this VLAN.</p> <p>If one physical network port (that is, a VLAN trunk) will handle multiple VLANs, create multiple VLAN subinterfaces on that port, one for each VLAN ID that will be received.</p> <p>If multiple different physical network ports will handle the same VLANs, on each of the ports, create VLAN subinterfaces that have the same VLAN IDs.</p> <p>The valid range is between 1 and 4094. The value you specify must match the VLAN ID added by the IEEE 802.1q-compliant router or switch connected to the VLAN subinterface.</p>
Interface	Physical interface associated with the VLAN; for example, port2.
Secondary IP List	
IP Address	<p>Secondary IP addresses can be used when you deploy the system so that it belongs to multiple logical subnets. If you assign multiple IP addresses to an interface, you must assign them static addresses.</p> <p>To add secondary IP addresses, enable the feature and save the configuration. After you have saved it the first time, you can edit it to add secondary IP addresses and enable inbound traffic to that address. For each address, specify an IP address using the CIDR-formatted subnet mask, separated by a forward slash (/), such as 192.0.2.5/24.</p>
Allow Access	Select the services that are allowed to send inbound traffic.
HA Node IP List	
IP Address	<p>You use the HA node IP list configuration in an HA active-active deployment. On each HA cluster node, add an HA node IP list that includes an entry for each cluster node. When the appliance is in standalone mode, it uses the physical port IP address; when it is in HA mode, it uses the HA node IP list address.</p> <p>For each address, specify an IP address using the CIDR-formatted subnet mask, separated by a forward slash (/), such as 192.0.2.5/24.</p>
Node ID	ID of the corresponding node.
Allow Access	Select the services that are allowed to send inbound traffic.

In an HA active-active deployment, if an interface uses secondary IP addresses, you must use the CLI to enable the HA node secondary IP address list, and then configure the list:



```
FADC # config system interface
FADC (interface) # edit port3
FADC (port3) # set ha-node-secondary-ip enable
FADC (port3) # config ha-node-secondary-ip-list
FADC (ha-node-second~r) # edit 1
Add new entry '1' for node 2221
FADC (1) # set ip 192.168.1.100
FADC (1) # set allowaccess https http ping snmp ssh
FADC (1) # end
FADC (port3) # end
```

To configure a physical interface using the CLI:

```
config system interface
edit <port_name>
set ip <ip&netmask>
set allowaccess {http https ping snmp ssh
telnet}
end
```

To configure an aggregate interface using the CLI:



```
config system interface
edit <specified_name>
set type agg
set aggregate-mode {802.3ad | balance-alb | balance-
rr | balance-tlb | balance-xor | broadcast}
set aggregate-algorithm {layer2 | layer2_3 | layer3_
4}
set member <port_name> <port_name>
set ip <ip&netmask>
end
```

To configure a VLAN interface using the CLI:

```
config system interface
edit <specified_name>
set type vlan
set vlanid <number>
set interface <port_name>
set ip <ip&netmask>
end
```

Configuring static routes

Network systems maintain route tables to determine where to forward TCP/IP packets. Routes for outbound traffic are chosen according to the following priorities:

- Link local routes—Self-traffic uses link local routes.
- LLB Link Policy route—Configured policy routes have priority over default routes.
- Policy route—Configured policy routes have priority over default routes.
- Static route / ISP route / OSPF route—Priority is based on the distance metric. By default, distance for static routes is 10, for ISP is 20, for OSPF is 110, for EBGp is 20, and IBGP is 200. The distance metric is configurable for static routes and OSPF routes, but not ISP routes.
- Default LLB Link Policy route—Default routes have lower priority than configured routes.
- Default static route / OSPF route—Default routes have lower priority than configured routes.

The system evaluates content route rules first, then policy routes, then static routes. The packets are routed to the first route that matches. The static route table, therefore, is the one that must include a “default route” to be used when no more specific route has been determined.

Static routes specify the IP address of a next-hop router that is reachable from that network interface. Routers are aware of which IP addresses are reachable through various network pathways, and can forward those packets along pathways capable of reaching the packets' ultimate destinations. The FortiADC system itself does not need to know the full route, as long as the routers can pass along the packet.

You must configure at least one static route that points to a router, often a router that is the gateway to the Internet. You might often classified as a path-vector protocol and sometimes as a distance-vector routing protocol, BGP exchanges routing and reachability information among autonomous systems over the Internet. You need to configure multiple static routes if you have multiple gateway routers, redundant ISP links, or other special routing cases.

Before you begin:

- You must have Read-Write permission for System settings.

To configure a static route:

1. Go to Networking > Routing.
The configuration page displays the Static tab.
2. Click **Add** to display the configuration editor.
3. Complete the configuration as described in [Table 100](#).
4. Save the configuration.

Table 100: Static route configuration

Settings	Guidelines
Destination	<p>Address/mask notation to match the destination IP in the packet header.</p> <p>It is a best practice to include a default route. If there is no other, more specific static route defined for a packet's destination IP address, a default route will match the packet, and pass it to a gateway router so that any packet can reach its destination. If you do not define a default route, and if there is a gap in your routes where no route matches a packet's destination IP address, packets passing through the FortiADC towards those IP addresses will, in effect, be null routed. While this can help to ensure that unintentional traffic cannot leave your FortiADC and therefore can be a type of security measure, the result is that you must modify your routes every time that a new valid destination is added to your network. Otherwise, it will be unreachable. A default route ensures that this kind of locally-caused "destination unreachable" problem does not occur. Specify 0.0.0.0/0 or ::/0 to set a default route for all packets.</p>
Gateway	<p>Specify the IP address of the next-hop router where the FortiADC system will forward packets for this static route. This router must know how to route packets to the destination IP addresses that you have specified, or forward packets to another router with this information. For a direct Internet connection, this will be the router that forwards traffic towards the Internet, and could belong to your ISP. The gateway must be in the same subnet as the interface used to reach it.</p>
Distance	<p>The default administrative distance is 10, which makes it preferred to OSPF routes that have a default of 110. We recommend you do not change these settings unless your deployment has exceptional requirements.</p>

To configure a static route using the CLI:



```
config router static
edit 1
set destination <ip address/netmask>
set gateway <ip address>
set distance <value>
end
```

Configuring policy routes

Network systems maintain route tables to determine where to forward TCP/IP packets. Policy routes set the gateway for traffic with a source and destination that match the policy.

Routes for outbound traffic are chosen according to the following priorities:

1. Link local routes—Self-traffic uses link local routes.
2. LLB Link Policy route—Configured policy routes have priority over default routes.
3. Policy route—Configured policy routes have priority over default routes.
4. Static route / ISP route / OSPF route—Priority is based on the distance metric. By default, distance for static routes is 10, for ISP routes is 20, and for OSPF routes is 110. The distance metric is configurable for static routes and OSPF routes, but not ISP routes.

5. Default LLB Link Policy route—Default routes have lower priority than configured routes.
6. Default static route / OSPF route—Default routes have lower priority than configured routes.

The system evaluates policy routes, then static routes. The packets are routed to the first route that matches. The policy route table, therefore, need not include a “default route” for packets that do not match your policy because those packets can be forwarded to the default route set in the static route table.

Most policy route settings are optional, so a matching route might not provide enough information to forward the packet. In that case, the FortiADC appliance may refer to the routing table in an attempt to match the information in the packet header with a route in the routing table. For example, if the destination address is the only match criteria in the policy route, the FortiADC appliance looks up the IP address of the next-hop router in its routing table. This situation could occur when interfaces are dynamic (such as DHCP or PPPoE) and you do not want or are unable to specify a static IP address of the next-hop router.

Before you begin:

- You must have Read-Write permission for System settings.

To configure a policy route:

1. Go to Networking > Routing.
2. Click the **Policy** tab.
3. Click **Add** to display the configuration editor.
4. Complete the configuration as described in [Table 101](#).
5. Save the configuration.

Table 101: Policy route configuration

Settings	Guidelines
Source	Address/mask notation to match the source IP in the packet header. To match any value, either leave it blank or enter 0.0.0.0/32.
Destination	Address/mask notation to match the destination IP in the packet header. To match any value, leave it blank or enter 0.0.0.0/32.
Gateway	IP address of the next-hop router where the FortiADC system will forward packets for this policy route. This router must know how to route packets to the destination subnet, or forward packets to another router with this information.

Chapter 12: System Management

This chapter includes the following topics:

- [Configuring basic system settings](#)
- [Configuring system time](#)
- [Backing up and restoring the configuration](#)
- [Updating firmware](#)
- [Configuring an SMTP mail server](#)
- [Configuring FortiGuard service settings](#)
- [Pushing/pulling configurations](#)
- [Configuring SNMP](#)
- [Manage and validate certificates](#)
- [HSM Integration](#)
- [Rebooting, resetting, and shutting down the system](#)

Configuring basic system settings

The basic system settings page includes configuration options for the following settings and features:

- Hostname
- Web UI language
- Management service ports
- DNS
- Virtual domain

Before you begin:

- You must have Read-Write permission for System settings.

To configure basic system settings:

1. Click System > Settings.
The configuration page displays the Basic tab.
2. Complete the configuration as described in [Table 102](#).
3. Save the configuration.

Table 102: Basic settings configuration

Settings	Guidelines
Hostname	<p>You can configure a hostname to facilitate system management. If you use SNMP, for example, the SNMP system name is derived from the configured hostname. The hostname can be up to 35 characters in length. It can include US-ASCII letters, numbers, hyphens, and under-scores, but not spaces and special characters.</p> <p>The System Information widget and the <code>get system status</code> CLI command display the full hostname. If the hostname is longer than 16 characters, the name is truncated and ends with a tilde (~) to indicate that additional characters exist, but are not displayed.</p>
Language	English or Simplified Chinese.
Idle Timeout	Log out an idle administrator session. The default is 30 minutes.
HTTP Port	Specify the port for the HTTP service. Usually, HTTP uses port 80.
HTTPS Port	Specify the port for the HTTPS service. Usually, HTTPS uses port 443.
Telnet Port	Specify the port for the Telnet service. Usually, Telnet uses port 25.
SSH Port	Specify the port for the SSH service. Usually, SSH uses port 22.
Primary DNS	<p>The system must be able to contact DNS servers to resolve IP addresses and fully qualified domain names. Your Internet service provider (ISP) might supply IP addresses of DNS servers, or you might want to use the IP addresses of your own DNS servers. You must provide unicast, non-local addresses for your DNS servers. Localhost and broadcast addresses are not accepted.</p> <p>Incorrect DNS settings or unreliable DNS connectivity can cause issues with other features, such as FortiGuard services and NTP system time.</p>
Secondary DNS	IPv4/IPv6 address of the secondary DNS server for your local network.
Virtual Domain	Enables the virtual domain feature. Before you enable it, make sure you understand how the system implements virtual domains. See Chapter 15: Virtual Domains .
Config Sync Enable	Enable/disable the configuration synchronization feature. This feature is related to Pushing/pulling configurations , not HA synchronization. Disabled by default.

Configuring system time

The system time must be accurate for many features to work, including scheduling, logging, and SSL/TLS-related features.

We recommend that you use Network Time Protocol (NTP) to maintain the system time. As an alternative when NTP is not available or is impractical, you can set the system time manually.

You can change the system time with the web UI or the CLI.

Before you begin:

- You must have Read-Write permission for System settings.

To configure the system time:

1. Go to System > Settings.
2. Click the **Maintenance** tab.
3. Complete the configuration as described in [Table 103](#).
4. Save your changes.

Table 103: System time configuration

Setting	Guidelines
System Time	Displays the system time. You can use NTP to set the system time, or use the controls to set the system time manually. Specify time in HH:MM:SS format.
Daylight Saving Time	Enable if you want the system to adjust its own clock when its time zone changes between daylight saving time (DST) and standard time.
Time Zone	Select the time zone where the appliance is located.
NTP	
NTP	Select to use NTP.
NTP Server	Specify a space-separated list of IP addresses or FQDNs for an NTP server or pool, such as <code>pool.ntp.org</code> . To find an NTP server, go to http://www.ntp.org .
Synchronizing Interval	Specify how often the system synchronizes its time with the NTP server. The default is 60 minutes. The valid range is 1-1440.

To configure NTP using the CLI:

```
config system time ntp
  set ntpsync enable
  set ntpserver {<server_fqdn> | <server_ipv4>}
  set syncinterval <minutes_int>
end
```



To configure the system time manually:

```
config system time ntp
  set ntpsync disable
end
config system time manual
  set zone <timezone_index>
  set daylight-saving-time {enable|disable}
end
execute date <MM/DD/YY> <HH:MM:SS>
```


Configuring an SMTP mail server

You can configure an SMTP email server if you want to send alerts by email. See [Configuring alert email settings](#) for information on alerts.

Before you begin:

- You must have Read-Write permission for System settings.

To configure SMTP:

1. Go to System > Settings.
2. Click the **Services** tab.
3. Complete the configuration as described in [Table 104](#).
4. Save the configuration.

Table 104: SMTP configuration

Settings	Guidelines
Address	IP address or FQDN of an SMTP server (such as FortiMail) or email server that the appliance can connect to in order to send alerts and/or generated reports.
Port	Listening port number of the server. Usually, SMTP is 25.
Authentication	Enable if the SMTP server requires authentication.
Security	STARTTLS is an extension to plain text communication protocols. It enables a plain text connection to be upgraded to an encrypted (TLS or SSL) connection instead of using a separate port for encrypted communication. Specify this option if you have implemented STARTTLS for your mailserver; otherwise, select none .
Username	Username for authentication to the SMTP server.
Password	Password for authentication to the SMTP server.

Configuring FortiGuard service settings

FortiGuard periodically updates the WAF Signature Database, IP Reputation Database, and Geo IP Database. You can go to the FortiGuard website to download the update packages that you can upload to FortiADC, or you can schedule automatic updates.

Before you begin:

- If you want to perform a manual update, you must download the update file from the FortiGuard website.

You must have Read-Write permission for System settings.

To configure FortiGuard service settings:

1. Go to System > Settings.
2. Click the **FortiGuard** tab.
3. Complete the configuration as described in [Table 105](#).
4. Save the configuration.

Table 105: FortiGuard service configuration

Settings	Guidelines
Support Contract	
Registration and license	<p>Review your registration and license information. If you need to change registration or if your license is about to expire, click Login Now to open the login page for the Fortinet Service & Support website in a new browser window.</p> <p>Note: If your license is invalid, FortiGuard does not send updates to the FortiADC. The functionality on FortiADC remains intact and useful, but it is out-of-date.</p>
FortiGuard Services	
WAF Signature Database	Review the version information. To perform a manual update, click Update to display controls that enable you to select and upload the update file.
IP Reputation Database	Review the version information. To perform a manual update, click Update to display controls that enable you to select and upload the update file.
Geo IP Database	Review the version information. To perform a manual update, click Update to display controls that enable you to select and upload the update file.
Configuration	
Scheduled Update Status	Enable updates.
Scheduled Update Frequency	<ul style="list-style-type: none"> • Every—Schedule periodic updates. Specify the time interval to perform updates. • Daily—Schedule daily updates. Specify the time of day to perform the update. • Weekly—Schedule weekly updates. Specify the day and time to perform the update.
Scheduled Update Day	Sunday, Monday, Tuesday, Wednesday, Thursday, Friday, Saturday.
Scheduled Update Time	HH:MM.HH is 0-23. MM is 0, 15, 30, or 45.
Override Server	If you are unable to make connections to the standard FortiGuard server, enable connection to the override server address given to you by Fortinet Service & Support.
Override Server Address	Override server IP address.
Web Filter Configure	

Settings	Guidelines
Cache Status	Enable/disable caching of the categorical lists of websites. FortiGuard maintains massive lists of web sites classified into categories so that you can enforce categorical decisions in your rules, like "do not do SSL forward proxy for sites belonging to the Personal Privacy category."
Cache TTL	Specify cache expiration. The default is 3600. The valid range is 10 to 86,400. When the cache expires, FortiADC initiates an update from FortiGuard.
FDS Port	Specify the port to receive updates. The default is 53. An alternative is 8888.

Pushing/pulling configurations

You can use the sync list configuration page to push or pull sets of configuration objects to or from a target FortiADC appliance. The push/pull operation is a manual operation. It is not repeated automatically.

Before you begin:

- Configuration synchronization must be enabled on the appliances. Go to System > Settings > Basic.
- You must plan for the impact the configuration push/pull has on the target deployment.
- You must have Read-Write permission for System settings.

To push or pull a configuration:

1. Click System > Settings.
2. Click the **Sync List** tab.
3. Click **Add** and complete the configuration as described in [Table 106](#).
After you have saved the configuration, it is added to the configuration table.
4. To execute the push/pull operation, select the configuration from the table, select **From** or **To**, and click **Sync**.
5. Check the Status column in the table to see the result of the push/pull operation.
6. Log into the target appliance and check the configuration logs (Log & Report > Log Browsing > Event Log > Configuration). Notice the log entries for each configuration change resulting from the push/pull operation.

Table 106: Sync List configuration

Settings	Guidelines
Name	Configuration name. Valid characters are A-Z, a-z, 0-9, _, and -. No spaces. After you initially save the configuration, you cannot edit the name.
Server IP	IP address of the remote appliance.
Password	Password for the admin account on the remote appliance.

Settings	Guidelines
Type	<ul style="list-style-type: none"> • System—Includes <code>config config</code>, <code>config system</code> (except <code>config system mailserver</code>), <code>config user</code>, and <code>config vdom</code> commands. • Route—Includes <code>config router</code> commands. • LB—Includes <code>config load-balance</code> commands. • FW—Includes <code>config firewall</code> commands. • Log—Includes <code>config log</code> commands and <code>config system mailserver</code>. • LLB—Includes <code>config link-load-balance</code> commands. • GDS—Includes <code>config global-load-balance</code> and <code>config global-dns-server</code> commands. • Security—Includes <code>config security waf</code> commands.
Comment	A string to describe the purpose of the configuration, to help you and other administrators more easily identify its use. Put phrases in quotes. For example: "SLB and GLB settings to Data Center East".

Backing up and restoring the configuration

You use the backup procedure to save a copy of the configuration. A simple backup file is a text file. A full backup is a tar file.

The backup feature has a few basic uses:

- Saving the configuration as CLI commands that a co-worker or Fortinet support can use to help you resolve issues with misconfiguration.
- Restoring the system to a known functional configuration.
- Creating a template configuration you can edit and then load into another system using the restore procedure.

A complete configuration backup is a tar file that includes the complete configuration files, plus any files you have imported, including error page files, script files, and ISP address book files.

In the event that FortiADC experiences hardware failure, being able to restore the entire backup configuration minimizes the time to reconfigure a replacement.

Configuration backups do *not* include data such as logs and reports.



Back up files can include sensitive information, such as HTTPS certificate private keys. We strongly recommend that you password-encrypt backup files and store them in a secure location.

Before you begin:

- If you are restoring a configuration, you must know its management interface configuration in order to access the web UI after the restore procedure is completed. Open the configuration file and make note of the IP address and

network requirements for the management interface (port1). You also must know the administrator username and password.

- You must have Read-Write permission for System settings.

To backup or restore the system configuration:

1. Click System > Settings.
2. Click the **Backup & Restore** tab.
3. Complete the actions described in [Table 107](#).

Table 107: Backup and restore configuration

Actions	Guidelines
Backup	
Back Up	Select this option to back up the configuration. This backup is a text file.
Local PC	Back up to the local PC.
ADC	Back up directly to the FortiADC device.
Entire Configuration	Select this option to include error page files, script files, and ISP address book files. This backup is a tar file.
Restore	
Restore (option)	Select this option to restore a previous configuration. This restore file must be a text file.
Local PC	Restore from the local PC.
ADC	Restore directly from your FortiADC device
File	Click Choose File to browse for the file to restore. Note: The option applies to restore operations from the local PC only.
Restore (button)	<p>Click the Restore button to start the restore procedure. Your web browser uploads the configuration file and the system restarts with the new configuration.</p> <p>Time required to restore varies by the size of the file and the speed of your network connection. Your web UI session is terminated when the system restarts. To continue using the web UI, refresh the web page and log in again.</p> <p>If the restored system has a different management interface configuration than the previous configuration, you must access the web UI using the new management interface IP address.</p>

Updating firmware

This topic includes the following information:

- [Upgrade considerations](#)
- [Updating firmware using the web UI](#)
- [Updating firmware using the CLI](#)

Upgrade considerations

The following considerations help you determine whether to follow a standard or non-standard upgrade procedure:

- HA—Updating firmware on an HA cluster requires some additions to the usual steps for a standalone appliance. For details, see [Updating firmware for an HA cluster](#).
- Re-imaging—If you are installing a firmware version that requires a different size of system partition, you might be required to re-image the boot device. Consult the release notes. In that case, do *not* install the firmware using this procedure. Instead, see [Restoring firmware \(“clean install”\)](#).
- Downgrades—If you are downgrading the firmware to a previous version, and the settings are not fully backwards compatible, the system might remove incompatible settings or use the default values for that version of the firmware. You might need to reconfigure some settings.

Important: Read the release notes for release-specific upgrade considerations.

Updating firmware using the web UI

[Figure 51](#) shows the user interface for managing firmware (either upgrades or downgrades). Firmware can be loaded on two disk partitions: the active partition and the alternate partition. The upgrade procedure:

- Updates the firmware on the inactive partition and then makes it the active partition.
- Copies the firmware on the active partition, upgrades it, and installs it in place of the configuration on the inactive partition.

For example, if partition 1 is active, and you perform the upgrade procedure:

- Partition 2 is upgraded and becomes the active partition; partition 1 becomes the alternate partition.
- The configuration on partition 1 remains in place; it is copied, upgraded, and installed in place of the configuration on partition 2.

The reason for this is to preserve the working system state in the event upgrade fails or is aborted.

Figure 51: User interface for managing firmware updates

The screenshot displays the FortiADC web interface under the 'Maintenance' tab. The left sidebar contains navigation options like Dashboard, System, Settings, and Shared Resources. The main content area is divided into three sections:

- Time:** Contains settings for System Time (2016-10-11 17:59:33), Daylight Saving Time (Enabled), Time Zone ((GMT-7:00)Pacific Time(US&Canada)), and NTP (Disabled). Buttons for 'Save' and 'Refresh' are present.
- Firmware:** Features a table with two partitions. Partition 1 is active (green checkmark) and Partition 2 is alternate (red X). The firmware version for both is FA-VMX-4.06.00-FW-build0660. A 'Boot Alternate Firmware' button is located below the table.
- Upgrade:** Includes a section for 'HA Sync' (Disabled) and a 'File' upload area with a 'Choose File' button and the text 'No file chosen. Select a file to upload.'

Before you begin:


- Download the firmware file from the Fortinet Customer Service & Support website: <https://support.fortinet.com/>
- Read the release notes for the version you plan to install.
- Back up your configuration before beginning this procedure. Reverting to an earlier firmware version could reset settings that are not compatible with the new firmware.
- You must have super user permission (user **admin**) to upgrade firmware.

To boot the firmware on the alternate partition:

- Click **Boot Alternate Firmware**.

The system reboots, the alternate becomes the active firmware, and the active becomes the alternate firmware.

To update firmware:

1. Go to System > Settings.
2. Click the **Maintenance** tab.
3. Scroll to the Upgrade section.
4. Click **Choose File** to locate and select the file.
5. Click  to upload the firmware and reboot.

The system replaces the firmware on the alternate partition and reboots. The alternate (upgraded) partition becomes the active, and the active becomes the alternate.

When you update software, you are also updating the web UI. To ensure the web UI displays the updated pages correctly:



- Clear your browser cache.
- Refresh the page.

In most environments, press Ctrl-F5 to force the browser to get a new copy of the content from the web application. See the Wikipedia article on browser caching issues for a summary of tips for many environments:

https://en.wikipedia.org/wiki/Wikipedia:Bypass_your_cache.

Updating firmware using the CLI

The CLI upgrade procedure replaces the firmware on the alternate partition and reboots. The alternate (upgraded) partition becomes the active, and the active becomes the alternate.

Note: The CLI does not have an equivalent of the web UI **Boot Alternative Firmware** command.

Before you begin:

- Read the release notes for the version you plan to install. If information in the release notes is different from this documentation, follow the instructions in the release notes.
- You must be able to use TFTP to transfer the firmware file to the FortiADC. Download and install a TFTP server, like `tftpd` ([Windows](#), [Mac OS X](#), or [Linux](#)), on a server on the same subnet as the FortiADC.
- Download the firmware file from the Fortinet Customer Service & Support website: <https://support.fortinet.com/>
- Copy the firmware image file to the root directory of the TFTP server.
- Back up your configuration before beginning this procedure.
- You must have super user permission (user **admin**) to upgrade firmware.



TFTP is not secure, and it does not support authentication. You should run it only on trusted administrator-only networks, and never on computers directly connected to the Internet. Turn off `tftpd` off immediately after completing this procedure.

To install firmware via the CLI:

1. Connect your management computer to the FortiADC console port using an RJ-45-to-DB-9 serial cable or a null-modem cable.
2. Initiate a connection to the CLI and log in as the user **admin**.
3. Use an Ethernet cable to connect FortiADC port1 to the TFTP server directly, or connect it to the same subnet as the TFTP server.
4. If necessary, start the TFTP server.
5. Use the following command to transfer the firmware image to the FortiADC system:

```
execute restore image tftp <filename> <tftp_ipv4>
```

The following example shows an upgrade:

```
FortiADC-VM # execute restore image tftp FAD_VM-v400-build0308-FORTINET.out 192.0.2.1
This operation will replace the current firmware version!
Do you want to continue? (y/n)y
```



```

Connect to tftp server 192.0.2.1 ...
Please wait...
#####
Get image from tftp server OK.
Check image trailer OK.
Check image OK.
FortiADC-VM #

```

The following example shows a downgrade:

```

FortiADC-VM # execute restore image tftp FAD_VM-v400-build0307-FORTINET.out 192.0.2.1
This operation will replace the current firmware version!
Do you want to continue? (y/n)y
Connect to tftp server 192.0.2.1 ...
Please wait...
#####
Get image from tftp server OK.
Check image trailer OK.
This operation will downgrade the current firmware version!
Do you want to continue? (y/n)y
FortiADC-VM #

```

6. To verify the upgrade, display the system version number:

```

FortiADC-VM # get system status
Version: FortiADC-VM v4.2.0,build0307,150209
VM Registration: Valid: License has been successfully authenticated with registration
servers.
VM License File: License file and resources are valid.
VM Resources: 1 CPU/1 allowed, 1620 MB RAM/2048 MB allowed, 23 GB Disk/1024 GB allowed
...

```



If the download fails after the integrity check with the error message `invalid compressed format (err=1)`, but the firmware matches the integrity checksum on the Fortinet Customer Service & Support website, try a different TFTP server.

Rebooting, resetting, and shutting down the system

The following items have the indicated usage:

- **Reboot**—Reboots the operating system.
- **Reset**—Resets the configuration to the default factory values.
- **Shut Down**—Shuts down the system. When the system is shut down, it is unavailable to forward traffic.



Do not unplug or switch off the FortiADC appliance without first shutting down the operating system. The shutdown process enables the system to finish writing any buffered data, and to correctly spin down and park the hard disks. Failure to do so could cause data loss and hardware problems.

To reboot the system:

Do one of the following:

- Go to the dashboard, and in the System Information widget, click **Reboot**.
- From the CLI console, enter the following command:

```
execute reboot
```

To perform a factory reset:

Do one of the following:

- Go to the dashboard, and in the System Information widget, click **Reset**.
- From the CLI console, enter the following command:

```
execute factoryreset
```

To power off the system:

To shut down the system:

- Go to the dashboard, and in the System Information widget, click **Shut Down**.
- From the CLI console, enter the following command:

```
execute shutdown
```

The system does not emit disk activity noise when shutdown is complete.

To completely power off:


- For hardware appliances, press the power button if there is one. Power supplies and switches vary by hardware model. On some, you press the power button; on others, you flip the switch to either the off (O) or on (I) position.
- For FortiADC-VM, power off the virtual machine.

Create a traffic group

A *traffic group* is a set of VRIDs. Each VRID keeps talking with its peers using 'hello' packets via its heartbeat interface so that each VRID can be aware of its peers (master or slave) operating state and perform a VRRP fail-over in case the master node fails. The different VRIDs have no relationship with each other.

In the example below, both VRID1 and VRID2 use Device1 as the master. When Port2 on Device1 fails, all traffic between the client and the server can't pass through the device

Figure 52: Traffic group



To solve this problem, you can create a traffic group and add both VRID1 and VRID2 as its members, and set the rule that the whole traffic group to fail over to the success device when either VRID fails. In this case, if Device1's Port2 fails, the whole traffic group will fail over to Device2.

Using the VRID concept, FortiADC allows you to add objects with floating IP address, such as interface, virtual server, IP pool, and SNA T pool, etc. to a traffic-group. With this configuration, it will trigger the whole traffic group to switch over when a resource fails.

Normally, the number of traffic groups should be same as the number of devices in an HA group for HA active-active configurations. FortiADC comes with a predefined traffic group named 'default'. You can use this default traffic group if you only need an HA active-passive deployment. Otherwise, you must configure your own traffic groups before making HA active-active configurations, using the instructions discussed in the following paragraphs.

Create a traffic group via the command line interface

Use the following commands to create a new traffic group:

```
config system traffic-group
  edit traffic-group-1
    set preempt enable
    set network-failover enable
    set failover-order 1 3 5
  next
end
```

Note: The failover sequence must be configured according to the order of node IDs. This means that if a node is dead, the next node in queue will take over handling the traffic. If you want to decide when a node should retake the traffic over from power-down to start-up, you **MUST** enable the Preempt option.

Create a traffic group from the Web GUI

Use the following steps to configure a traffic group from FortiADC's web interface:

1. Click System > Traffic Group.
2. Click **Add** to open the Traffic Group dialog.
3. Make the desired entries or selections as described in the table below.
4. Click **Save** when done.

Table 108: Traffic-group parameters

Parameter	Description
Traffic Group Name	Specify a unique name for the traffic group.
Preempt	Disabled by default. If enabled, the node will retake control of traffic from power-down to start-up.
Network Failover	Disabled by default. If enabled, the node will perform failover when the node's remote IP monitor detects failure of the node itself.
Failover Order	Follow the hint onscreen to set the failover sequence among the ports.

Create administrator users

We recommend that only network administrators—and if possible, only a single person—use the **admin** account. You can configure accounts that provision different scopes of access. For example, you can create an account for a security auditor who must only be able to view the configuration and logs, but *not* change them.

Before you begin:

- If you want to use RADIUS or LDAP authentication, you must have already have created the RADIUS server or LDAP server configuration.
- You must have Read-Write permission for System settings.

To create an administrator user account:

1. Go to System > Administrator.
2. Make sure the Admin tab is selected.
3. Click **Add** to display the configuration editor.
4. Complete the configuration as described in [Table 109](#).
5. Click **Save**.

Table 109: Administrator user configuration

Settings	Guidelines
Name	<p>Name of the administrator account, such as <code>admin1</code> or <code>admin@example.com</code>.</p> <p>Do not use spaces or special characters except the 'at' symbol (<code>@</code>). The maximum length is 35 characters.</p> <p>If you use LDAP or RADIUS, specify the LDAP or RADIUS username. This is the user name that the administrator must provide when logging in to the CLI or web UI. The users are authenticated against the associated LDAP or RADIUS server.</p> <p>After you initially save the configuration, you cannot edit the name.</p>

Settings	Guidelines
Trusted Hosts	<p>Source IP address and netmask from which the administrator is allowed to log in. For multiple addresses, separate each entry with a space. You can specify up to three trusted areas. They can be single hosts, subnets, or a mixture.</p> <p>Configuring trusted hosts hardens the security of the system. In addition to knowing the password, an administrator must connect only from the computer or subnets you specify.</p> <p>Trusted host definitions apply both to the web UI and to the CLI when accessed through Telnet, SSH, or the CLI console widget. Local console access is <i>not</i> affected by trusted hosts, as the local console is by definition not remote, and does not occur through the network.</p> <p>If ping is enabled, the address you specify here is also a source IP address to which the system will respond when it receives a ping or traceroute signal.</p> <p>To allow logins only from <i>one</i> computer, enter only its IP address and 32- or 128-bit netmask:</p> <pre>192.0.2.1/32</pre> <pre>2001:0db8:85a3::8a2e:0370:7334/128</pre> <p>To allow login attempts from any IP address (not recommended), enter:</p> <pre>0.0.0.0/0</pre> <p>Caution: If you restrict trusted hosts, do so for <i>all</i> administrator accounts. Failure to do so means that all accounts are still exposed to the risk of brute force login attacks. This is because if you leave even <i>one</i> administrator account unrestricted (i.e. 0.0.0.0/0), the system must allow login attempts on all network interfaces where remote administrative protocols are enabled, and wait until <i>after</i> a login attempt has been received in order to check that user name's trusted hosts list.</p> <p>Tip: If you allow login from the Internet, set a longer and more complex New Password, and enable only secure administrative access protocols. We also recommend that you restrict trusted hosts to IPs in your administrator's geographical area.</p> <p>Tip: For improved security, restrict all trusted host addresses to single IP addresses of computer(s) from which <i>only</i> this administrator will log in.</p>
Global Admin	<ul style="list-style-type: none"> • No —Default. If selected, the account can access the virtual domain specified in this configuration only. • Yes—If selected, the account can access all virtual domains.

Settings	Guidelines
Profile	<p>Select a user-defined or predefined profile. The predefined profile named super_admin_prof is a special access profile used by the admin account. However, selecting this access profile will <i>not</i> confer all permissions of the admin account. For example, the new administrator would not be able to reset lost administrator passwords.</p> <p>Note: This option does not appear for the admin administrator account, which by definition always uses the super_admin_prof access profile.</p>
Scope	Select a user-defined or predefined scope.
Authentication Server	<ul style="list-style-type: none"> Local—Use the local administrator authentication server. RADIUS—Use a RADIUS authentication server. Select the RADIUS server configuration. LDAP—Use an LDAP authentication server. Select the LDAP server configuration.
Password	Set a strong password for all administrator accounts. The password should be at least eight characters long, be sufficiently complex, and be changed regularly. To check the strength of your password, you can use a utility such as Microsoft's password strength meter .
Virtual Domain	Optional. If you have enabled the virtual domain feature, select the virtual domain that this administrator can view and manage.

Configure access profiles

Access profiles provision permissions to roles. The following permissions can be assigned:

- Read (view access)
- Read-Write (view, change, and execute access)
- No access

When an administrator has only read access to a feature, the administrator can access the web UI page for that feature, and can use the `get` and `show` CLI command for that feature, but cannot make changes to the configuration.

In larger companies where multiple administrators divide the share of work, access profiles often reflect the specific job that each administrator does ("role"), such as account creation or log auditing. Access profiles can limit each administrator account to their assigned role. This is sometimes called role-based access control (RBAC).

Table 110 lists the administrative areas that can be provisioned. If you provision read access, the role can view the web UI menu (or issue a CLI `get` command). If you provision read-write access, the role can save configuration changes (or issue a CLI `set` command).

For complete access to *all* commands and abilities, you must log in with the administrator account named **admin**.

Table 110: Areas of control in access profiles

Web UI Menus	CLI Commands
System	<code>config system</code> <code>diagnose hardware</code> <code>diagnose sniffer</code> <code>diagnose system</code> <code>execute date</code> <code>execute ping</code> <code>execute ping-options</code> <code>execute traceroute</code>
Router	<code>config router</code>
Server Load Balance	<code>config load-balance</code>
Link Load Balance	<code>config link-load-balance</code>
Global Load Balance	<code>config global-dns-server</code> <code>config global-load-balance</code>
Security	<code>config firewall</code> <code>config security waf</code>
Log & Report	<code>config log</code> <code>config report</code> <code>execute rebuild-db</code>
<p>* For each <code>config</code> command, there is an equivalent <code>get/show</code> command. The <code>config</code> commands require write permission. The <code>get/show</code> commands require read permission.</p>	

Before you begin:

- You must have Read-Write permission for System settings.

To configure administrator profiles:

- Click System > Administrator.
- Click the **Access Profile** tab.
- Click **Add** to display the configuration editor.
- Complete the configuration as described in [Table 111](#).
- Click **Save**.

Table 111: Access profile configuration

Settings	Guidelines
Name	Configuration name. Valid characters are A-Z, a-z, 0-9, _, and -. No spaces.

Settings	Guidelines
System	<p>Select one of the following:</p> <ul style="list-style-type: none"> • None—Do not provision access for the menu. • Read Only—Provision ready-only access. • Read-Write—Enable the role to make changes to the configuration.
Networking	<p>Select one of the following:</p> <ul style="list-style-type: none"> • None—Do not provision access for the menu. • Read Only—Provision ready-only access. • Read-Write—Enable the role to make changes to the configuration.
User	<p>Select one of the following:</p> <ul style="list-style-type: none"> • None—Do not provision access for the menu. • Read Only—Provision ready-only access. • Read-Write—Enable the role to make changes to the configuration.
Server Load Balance	<p>Select one of the following:</p> <ul style="list-style-type: none"> • None—Do not provision access for the menu. • Read Only—Provision ready-only access. • Read-Write—Enable the role to make changes to the configuration.
Link Load Balance	<p>Select one of the following:</p> <ul style="list-style-type: none"> • None—Do not provision access for the menu. • Read Only—Provision ready-only access. • Read-Write—Enable the role to make changes to the configuration.
Global Load Balance	<p>Select one of the following:</p> <ul style="list-style-type: none"> • None—Do not provision access for the menu. • Read Only—Provision ready-only access. • Read-Write—Enable the role to make changes to the configuration.
Security	<p>Select one of the following:</p> <ul style="list-style-type: none"> • None—Do not provision access for the menu. • Read Only—Provision ready-only access. • Read-Write—Enable the role to make changes to the configuration.
Log & Report	<p>Select one of the following:</p> <ul style="list-style-type: none"> • None—Do not provision access for the menu. • Read Only—Provision ready-only access. • Read-Write—Enable the role to make changes to the configuration.

Settings	Guidelines
Shared Resource	<p>For each category, set the permission:</p> <ul style="list-style-type: none"> • None—Do not provision access for the menu. • Read Only—Provision ready-only access. • Read-Write—Enable the role to make changes to the configuration.



The **super_admin_prof** access profile, a special access profile assigned to the **admin** account and required by it, appears in the list of access profiles. It exists by default and cannot be changed or deleted. The profile has permissions similar to the UNIX root account.

Enable password policies

A password policy is a set of rules designed to enhance computer security. A good password policy encourages users to create strong passwords and use them properly. For your network and data security and integrity, we strongly recommend the enforcement of strong password policies when using FortiADC.

To enable password policy:

1. Go to System > Administrator.
2. Select the Password Policy tab.
3. Check the **Enable** check box next to Password Policy..
4. Complete the configuration as described in [Table 112](#).
5. Click **Save**.

Table 112: Password policy configuration

Settings	Guidelines
Minimum Length	Specify the minimum length of password, which can contain 8 (default) to 32 characters in length.
Must Contain	<p>Select the options you want to apply:</p> <ul style="list-style-type: none"> • Upper Case Letter—If selected, passwords must contain upper-case letters. • Lower Case Letter—If selected, passwords must contain lower-case letters. • Number—If selected, passwords must contain numbers. • Non Alphanumeric —If selected, passwords must contain non-alphanumeric characters.

Configuring SNMP

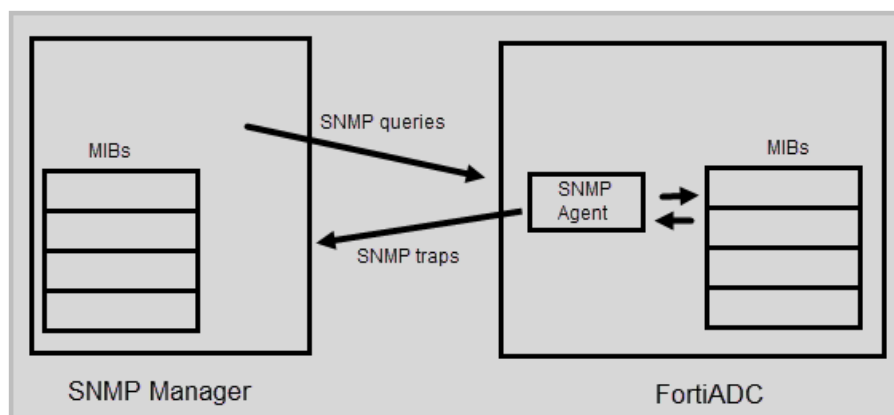
Many organizations use *SNMP* (simple network management protocol) to track the health of their systems. FortiADC supports SNMP v1, v2c, and v3.

SNMP depends on network devices that maintain standard management information bases (MIBs). *MIBs* describe the structure of the management data maintained on the device. Some MIB definitions are standard for all network devices, and some are vendor and product-family specific.

The FortiADC system runs an *SNMP agent* to communicate with the *SNMP manager*. The agent enables the system to respond to *SNMP queries* for system information and to send *SNMP traps* (alarms or event messages) to the SNMP manager.

Figure 53 illustrates the basic communication.

Figure 53: SNMP communication



With SNMP v1 and v2c managers, you configure *SNMP communities* to connect FortiADC and the SNMP manager. The SNMP Manager sends the community string along with all SNMP requests. If the community string is correct, the device responds with the requested information. If the community string is incorrect, the device simply discards the request and does not respond.



Fortinet strongly recommends that you do not add FortiADC to the community named `public`. This default name is well-known, and attackers that attempt to gain access to your network often try this name first.

With SNMPv3 managers, you configure *SNMP users* to connect FortiADC and the SNMP manager. Queries and traps include username/password authentication, along with an encryption key. FortiADC implements the user security model described in [RFC 3414](#).

Before you begin:

- On the SNMP manager, you must verify that the SNMP manager is a member of the community to which the FortiADC system belongs, and you must compile the necessary Fortinet-proprietary management information blocks (MIBs) and Fortinet-supported standard MIBs. For information on Fortinet MIBs, see [Appendix A: Fortinet MIBs](#).

- In the FortiADC interface settings, you must enable SNMP access on the network interface through which the SNMP manager connects.
- You must have Read-Write permission for System settings.

To configure SNMP system information:

1. Go to System > SNMP.
2. Click the System Information tab.
3. Complete the configuration as described in [Table 113](#).
4. Save the configuration.

Table 113: SNMP settings

Settings	Guidelines
SNMP Agent	Enable to activate the SNMP agent, so that the system can send traps and receive queries.
Description	A description or comment about the system, such as <code>dont-reboot</code> . The description can be up to 35 characters long, and can contain only letters (a-z, A-Z), numbers, hyphens (-) and underscores (_).
Contact	Contact information for the administrator or other person responsible for this system, such as a phone number (555-5555) or name (jdoe). The contact information can be up to 35 characters long, and can contain only letters (a-z, A-Z), numbers, hyphens (-) and underscores (_).
Location	Physical location of the appliance, such as <code>floor2</code> . The location can be up to 35 characters long, and can contain only letters (a-z, A-Z), numbers, hyphens (-) and underscores (_).

Download SNMP MIBs

FortiADC allows you to download full FortiADC and Fortinet Core MIB files, which provides more options for server load balance, global server load balance, link load balance, and firewall with SNMP traps.

To download an SNMP MIB file:

1. Click System > SNMP.
2. Click the System Information tab.
3. In the FortiADC SNMP MIB section, click **Download FortiADC MIB File** or **Download Fortinet Core MIB File**.
4. Follow the instructions onscreen to complete the download.

Configure SNMP threshold

To configure SNMP threshold:

1. Go to System > SNMP.
2. Click the Threshold tab.

3. Complete the configuration as described in [Table 114](#).
4. Save the configuration.

Table 114: SNMP threshold

Settings	Guidelines
CPU	<ul style="list-style-type: none">• Trigger—The default is 80% utilization.• Threshold—The default is 3, meaning the event is reported when the condition has been triggered 3 times in a short period.• Sample Period—The default is 600 seconds.• Sample Frequency—The default is 30 seconds.
Memory	<ul style="list-style-type: none">• Trigger—The default is 80% utilization.• Threshold—The default is 3, meaning the event is reported when the condition has been triggered 3 times in a short period.• Sample Period—The default is 600 seconds.• Sample Frequency—The default is 30 seconds.
Disk	<ul style="list-style-type: none">• Trigger—The default is 90% utilization.• Threshold—The default is 1, meaning the event is reported each time the condition is triggered.• Sample Period—The default is 7200 seconds.• Sample Frequency—The default is 3600 seconds.

Configure SNMP v1/v2

To configure SNMP v1/v2:

1. Go to System > SNMP.
2. Click the SNMPv1/v2 tab.
3. Complete the configuration as described in [Table 115](#).
4. Save the configuration.

Table 115: SNMP settings

Settings	Guidelines
Name	<p>Name of the SNMP community to which the FortiADC system and at least one SNMP manager belongs, such as <code>management</code>.</p> <p>You must configure the FortiADC system to belong to at least one SNMP community so that community's SNMP managers can query system information and receive SNMP traps.</p> <p>You can add up to three SNMP communities. Each community can have a different configuration for queries and traps, and the set of events that trigger a trap.</p> <p>You can also add the IP addresses of up to eight SNMP managers to each community to designate the destination of traps and which IP addresses are permitted to query the FortiADC system.</p>
Status	Select to enable the configuration.
Queries	<p>Port number on which the system listens for SNMP queries from the SNMP managers in this community. The default is 161.</p> <p>Enable queries for SNMP v1, SNMP v2c, or both.</p>
Traps	<p>Source (Local) port number and destination (Remote) port number for trap packets sent to SNMP managers in this community. The default is 162.</p> <p>Enable traps for SNMP v1, SNMP v2c, or both.</p>
Events	<p>Select to enable SNMP event reporting for the following thresholds:</p> <ul style="list-style-type: none"> • CPU—CPU usage has exceeded 80%. • Memory—Memory (RAM) usage has exceeded 80%. • Log disk usage—Disk space usage for the log partition or disk has exceeded 90%.
Host	<ul style="list-style-type: none"> • IP Address—Subnet address for the SNMP manager to receive traps and be permitted to query the FortiADC system. SNMP managers have read-only access. You can add up to 8 SNMP managers to each community. To allow any IP address using this SNMP community name to query the FortiADC system, enter <code>0.0.0.0/0</code>. For security best practice reasons, however, this is not recommended. • Host Type—Whether the host can send queries, receive traps, or any (both). <p>Caution: The system sends security-sensitive traps, which should be sent only over a trusted network, and only to administrative equipment.</p> <p>Note: If there are no other host IP entries, entering only <code>0.0.0.0/0</code> effectively disables traps because there is no specific destination for trap packets. If you do not want to disable traps, you must add at least one other entry that specifies the IP address of an SNMP manager.</p>



Test both traps and queries (assuming you have enabled both). Traps and queries typically occur on different port numbers, and therefore verifying one does not necessarily verify that the other is also functional.

To test queries, from your SNMP manager, query the FortiADC appliance.

To test traps, cause one of the events that should trigger a trap.

Configure SNMP v3

To configure SNMP v3:

1. Go to System > SNMP.
2. Click the SNMPv3 tab.
3. Complete the configuration as described in [Table 116](#).
4. Save the configuration.

Table 116: SNMP v3

Settings	Guidelines
Name	User name that the SNMP Manager uses to communicate with the SNMP Agent. After you initially save the configuration, you cannot edit the name.
Status	Enable/disable the configuration.
Security Level	<ul style="list-style-type: none"> • No Auth And No Privacy—Do not require authentication or encryption. • Auth But No Privacy—Authentication based on MD5 or SHA algorithms. Select an algorithm and specify a password. • Auth And Privacy—Authentication based on MD5 or SHA algorithms, and encryption based on AES or DES algorithms. Select an Auth Algorithm and specify an Auth Password; and select a Private Algorithm and specify a Private Password.
Queries	<p>Port number on which the system listens for SNMP queries from the SNMP managers for this user. The default is 161.</p> <p>Enable queries for SNMP v3.</p>
Traps	<p>Source (Local) port number and destination (Remote) port number for trap packets sent to SNMP managers for this user. The default is 162.</p> <p>Enable traps for SNMP v3.</p>

Settings	Guidelines
Events	<p>Select to enable SNMP event reporting for the following thresholds:</p> <ul style="list-style-type: none"> • CPU—CPU usage has exceeded 80%. • Memory—Memory (RAM) usage has exceeded 80%. • Log disk usage—Disk space usage for the log partition or disk has exceeded 90%. • System—Reserved for future use. • RAID—Reserved for future use. • HA—Reserved for future use. • Remote Storage—Reserved for future use. • IP Change—Reserved for future use.
Host	<ul style="list-style-type: none"> • Host Type—Whether the host can send queries, receive traps, or any (both). • IP Address—Subnet address for the SNMP manager to receive traps and be permitted to query the FortiADC system. SNMP managers have read-only access. You can add up to 8 SNMP managers to each community. To allow any IP address using this SNMP community name to query the FortiADC system, enter 0.0.0.0/0. For security best practice reasons, however, this is not recommended. <p>Caution: The system sends security-sensitive traps, which should be sent only over a trusted network, and only to administrative equipment.</p> <p>Note: If there are no other host IP entries, entering only 0.0.0.0/0 effectively disables traps because there is no specific destination for trap packets. If you do not want to disable traps, you must add at least one other entry that specifies the IP address of an SNMP manager.</p>



Test both traps and queries (assuming you have enabled both). Traps and queries typically occur on different port numbers, and therefore verifying one does not necessarily verify that the other is also functional.

To test queries, from your SNMP manager, query the FortiADC appliance.

To test traps, cause one of the events that should trigger a trap.

Manage and validate certificates

This section includes the following topics:

- [Overview](#)
- [Prerequisite tasks](#)
- [Manage certificates](#)
- [Validate certificates](#)

Overview

The FortiADC system is able to process the following two types of TLS/SSL traffic:

- System administration—Administrators connect to the web UI (HTTPS connections only). When you connect to the web UI, the system presents its own default “Factory” certificate. This certificate is used only for connections to the web UI. It cannot be removed. Do not use this certificate for server load balancing traffic.
- Server load balancing—Clients use SSL or TLS to connect to a virtual server. When you use FortiADC as a proxy for SSL operations normally performed on the backend real servers, you must import the X.509 v3 server certificates and private keys that the backend servers would ordinarily use, as well as any certificate authority (CA) or intermediate CA certificates that are used to complete the chain of trust between your clients and your servers.

The FortiADC system supports all of the TLS/SSL administration methods commonly used by HTTPS servers, including:

- **Server name indication (SNI)**—You can require clients to use the TLS extension to include the server hostname in the TLS client hello message. Then, the FortiADC system can select the appropriate local server certificate to present to the client.
- Local certificate store—A certificate store for the X.509 v3 server certificates and private keys that the backend servers would ordinarily use.
- Intermediate CAs store—A store for Intermediate CAs that the backend servers would ordinarily use to complete the chain of server certificates. HTTPS transactions use intermediate CAs when the server certificate is signed by an intermediate certificate authority (CA) rather than a root CA.
- Certificate Authorities (CAs) store—A store for CA certificates that the back-end servers would ordinarily use to verify the CA signature in client certificates or the signature of an OCSP Responder.
- OCSP—Use Online Certificate Status Protocol (OCSP) to obtain the revocation status of certificates.
- CRL—Use a Certificate Revocation List (CRL) to obtain the revocation status of certificates.
- Certificate validation policy—You can configure certificate validation policies that use OCSP or CRL. These policies can be associated with load balancing profiles.
- All digital certificates of RSA and ECDSA key types—whether they are local, remote, intermediate, or CA root certificates.
- Multiple CA, CRL, and OCSP configurations.
- Client certificate forwarding
- SNI forward ing
- Email alert of certificate expiration

Certificates and their domains

You can generate or import certificates in the global domain (i.e., FortiADC appliance) and individual VDOM domains (i.e., virtual machines). The visibility and use of certificates or certificate groups may vary, depending where (the domain) they are created. Below are the general guidelines regarding the availability and use of certificates or certificate groups.

- **Local Certificates/intermediate Certificates**—If generated or imported in a specific VDOM domain, they can be viewed and deleted in that VDOM only, but not visible in any other VDOM or the global domain; if generated or imported in the global domain, they can be viewed and downloaded by all VDOMS, but can be deleted only in the global domain.

- **Local Certificate Groups/Intermediate CA Groups**—If added in a specific VDOM domain, they can be viewed, edited, or referenced in that VDOM domain only, but not visible in any other VDOMs or the global domain; if added in the global domain, they can be visible to all VDOM domains, but can be edited only in the global domain.
- **CA/CRL/OCSP Signing Certificates**—If imported in a specific VDOM domain, they can be viewed or deleted only in that VDOM, but not visible in any other VDOM domain or the global domain; if imported in the global domain, they can be viewed or downloaded in all VDOM domains, but can be deleted only in the global domain.
- **Verify/CA Group/OCSP**—If added in a specific VDOM domain, they can be viewed or edited or referenced to in that VDOM domain only, but not visible in any other VDOM domain or the global domain; if added in the global domain, they can be viewed or referenced to in all VDOMs, but can be edited only in the global domain.

Prerequisite tasks

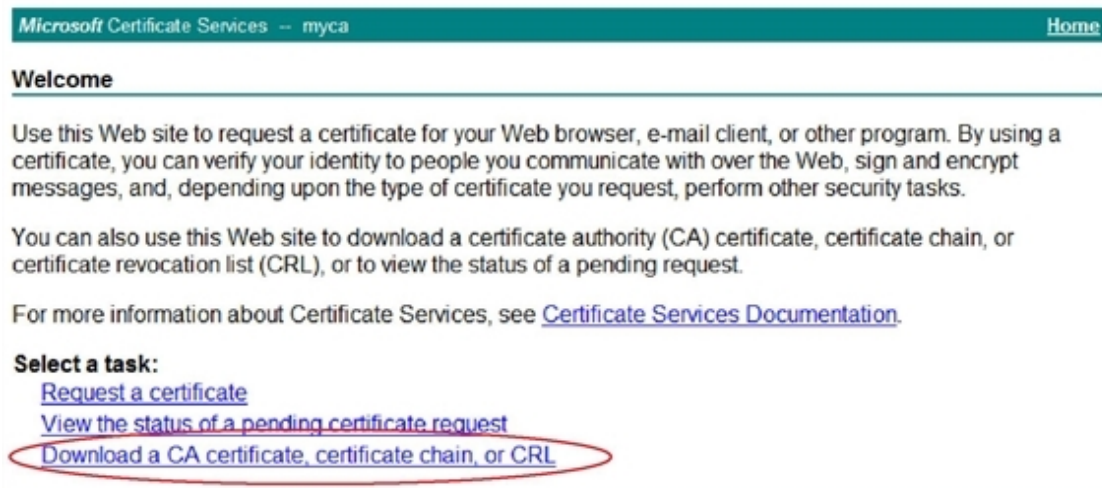
You must download the certificates from your backend servers so that you can import them into the FortiADC system.

This example shows how to download a CA certificate from Microsoft Windows 2003.

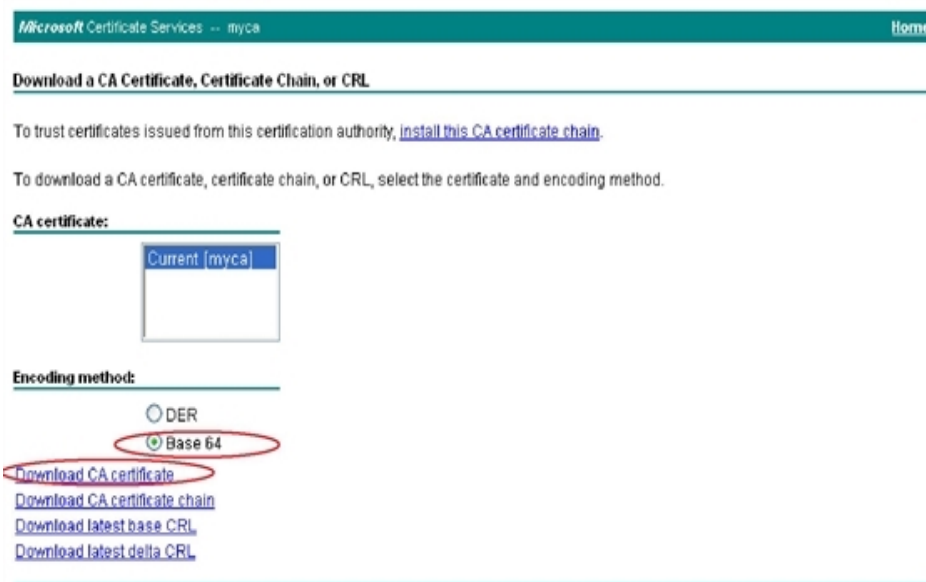
To download a CA certificate from Microsoft Windows 2003 Server:

1. Go to `https://<ca-server_ipv4>/certsrv/`.
where `<ca-server_ipv4>` is the IP address of your CA server.
2. Log in as Administrator. Other accounts may not have sufficient privileges.
The Microsoft Certificate Services home page appears. [Figure 54](#) is an example of this page.

Figure 54: Welcome page



3. Click the **Download CA certificate, certificate chain, or CRL** link to display the Download a CA Certificate, Certificate Chain, or CRL page. [Figure 55](#) is an example of this page.
4. From Encoding Method, select **Base64**.
5. Click **Download CA certificate**.

Figure 55: Download a CA Certificate, Certificate Chain, or CRL page

Manage certificates

This section discusses the following tasks you can perform on the System > Certificate > Manage Certificates page:

- Generating a certificate signing request
- Importing local CAs
- Importing intermediate CAs
- Creating an intermediate CA group
- Creating a local CA group

Generating a certificate signing request

Many commercial certificate authorities (CAs) provide websites where you can generate your own certificate signing request (CSR). A CSR is an unsigned certificate file that the CA will sign. When a CSR is generated, the associated private key that the appliance will use to sign and/or encrypt connections with clients is also generated.

If your CA does not provide this service, or if you have your own private CA such as a Linux server with OpenSSL, you can use FortiADC to generate a CSR and private key. This CSR can then be submitted for verification and signing by the CA.

Before you begin:

- You must have Read-Write permission for System settings.

To generate a certificate signing request:

1. Go to System > Certificate > Manage Certificates.
2. Click the **Local Certificate** tab.

3. Click **Generate** to display the configuration editor.
4. Complete the configuration as described in [Table 117](#).
5. Click **Save** when done.

The system creates a private and public key pair. The generated request includes the public key of the FortiADC appliance and information such as the IP address, domain name, or email address. The FortiADC appliance private key remains confidential on the FortiADC appliance. The Status column of the new CSR entry is **Pending**.

6. Select the row that corresponds to the certificate request.
7. Click **Download**.
Standard dialogs appear with buttons to save the file at a location you select. Your web browser downloads the certificate request (.csr) file.
8. Upload the certificate request to your CA.
After you submit the request to a CA, the CA will verify the information in the certificate, give it a serial number, an expiration date, and sign it with the public key of the CA.
9. If you are not using a commercial CA whose root certificate is already installed by default on web browsers, download your CA's root certificate, and then install it on all computers that will be connecting to your FortiADC appliance. Otherwise, those computers might not trust your new certificate.
10. After you've received the signed certificate from the CA, import the certificate into the FortiADC system.

Table 117: CSR configuration

Settings	Guidelines
Generate Certificate Signing Request	
Certification Name	Configuration name. Valid characters are A-Z, a-z, 0-9, _, and -. No spaces. The maximum length is 35 characters. Note: This is the name of the CSR file, not the host name/IP contained in the certificate's <code>Subject :</code> line.
Subject Information	

Settings	Guidelines
ID Type	<p>Select the type of identifier to use in the certificate to identify the virtual server:</p> <ul style="list-style-type: none"> • Host IP—The <i>static</i> public IP address of the FortiADC virtual server in the IP Address field. If the FortiADC appliance does not have a static public IP address, use the email or domain name options instead. Note: Do NOT use this option if your network has a dynamic public IP address. Your web browser will display the “Unable to verify certificate” or similar error message when your public IP address changes. • Domain Name—The fully qualified domain name (FQDN) of the FortiADC virtual server, such as <code>www.example.com</code>. This does not require that the IP address be static, and may be useful if, for example, your network has a dynamic public IP address and therefore clients connect to it via dynamic DNS. Do not include the protocol specification (<code>http://</code>) or any port number or path names. • E-Mail—The email address of the owner of the FortiADC virtual server. Use this if the virtual server does not require either a static IP address or a domain name. <p>Depending on your choice for ID Type, related options appear.</p>
IP Address	<p>Enter the static IP address of the FortiADC appliance, such as <code>10.0.0.1</code>. The IP address should be the one that is visible to clients. Usually, this should be its public IP address on the Internet, or a virtual IP that you use NAT to map to the appliance’s IP address on your private network.</p> <p>This option appears only if ID Type is Host IP.</p>
Domain Name	<p>Enter the FQDN of the FortiADC appliance, such as <code>www.example.com</code>. The domain name <i>must</i> resolve to the IP address of the FortiADC appliance or backend server according to the DNS server used by clients. (If it does not, the clients’ browsers will display a <code>Host name mismatch</code> or similar error message.)</p> <p>This option appears only if ID Type is Domain Name.</p>
Email	<p>Enter the email address of the owner of the FortiADC appliance, such as <code>admin@example.com</code>. This option appears only if ID Type is E-Mail.</p>
Distinguished Information	
Organization Unit	Name of organizational unit (OU), such as the name of your department. This is optional. To enter more than one OU name, click the + icon, and enter each OU separately in each field.
Organization	Legal name of your organization.
Locality (City)	City or town where the FortiADC appliance is located.
State/Province	State or province where the FortiADC appliance is located.
Country/Region	Country where the FortiADC appliance is located.
Email	E-mail address that may be used for contact purposes, such as <code>admin@example.com</code> .

Settings	Guidelines
Key Information	
Key Type	Select either of the following: <ul style="list-style-type: none"> • RSA • ECDSA
Key Size/ Curve Name	For RSA key, select one of the following key sizes: <ul style="list-style-type: none"> • 512 Bit • 1024 Bit • 1536 Bit • 2048 Bit • 4096 Bit. <p>Note: Larger keys use more computing resources, but provide better security.</p> <p>For ECDSA, select one of the following curve names:</p> <ul style="list-style-type: none"> • prime256v1 • secp384r1 • secp521r1
Enrollment Information	
Enrollment Method	<ul style="list-style-type: none"> • File-Based—You must manually download and submit the resulting certificate request file to a CA for signing. Once signed, upload the local certificate. <p>Online SCEP—The FortiADC appliance automatically uses HTTP to submit the request to the simple certificate enrollment protocol (SCEP) server of a CA, which will validate and sign the certificate. For this selection, two options appear. Enter the CA Server URL and the Challenge Password.</p>

Importing local certificates

You can import (upload) the following types of X.509 server certificates and private keys into the FortiADC system:

- Base64-encoded
- PKCS #12 RSA-encrypted

Before you begin:

- You must have Read-Write permission for System settings.
- You must have downloaded the certificate and key files and be able to browse to them
- so that you can upload them.

To import a local certificate:

1. Go to System > Certificate > Manage Certificates.
2. Click the **Local Certificate** tab.

3. Click **Import** to display the configuration editor.
4. Complete the configuration as described in [Table 118](#).
5. Click **Save** when done.

Table 118: Local certificate import configuration

Settings	Guidelines
Type	<p>Click the down arrow and select one of the following options from the drop-down menu:</p> <ul style="list-style-type: none"> • Local Certificate—Use this option only if you have a CA-signed certificate that was originated from a CSR generated in FortiADC . See Generating a certificate signing request on page 314. Note: It is important to make sure that the load-balancer (FortiADC appliance) you use to import a local certificate is the same appliance where the CSR was generated because it is where the key matching the certificate resides. The import operation will fail without the matching key on the same hardware system. • PKCS12 Certificate—Use this option only if you have a PKCS #12 password-encrypted certificate with its key in the same file. • Certificate—Use this option only if you have a certificate and its key in separate files. <p>Note: Additional fields are displayed depending on your selection.</p>
Local Certificate	
Certificate File	Browse for and upload the certificate file that you want to use.
PKCS12 Certificate	
Certificate Name	Specify the certificate name that can be referenced by other parts of the configuration, such as <code>www_example_com</code> . The maximum length is 35 characters. Do not use spaces or special characters.
Certificate File	Browse for and upload the certificate file that you want to use.
Password	Specify the password to encrypt the file in local storage.
Certificate	
Certificate Name	Specify the name that can be referenced by other parts of the configuration, such as <code>www_example_com</code> . The maximum length is 35 characters. Do not use spaces or special characters.
Certificate File	Browse for and upload the certificate file that you want to use.
Key File	Browse for and upload the corresponding key file.
Password	Specify the password to encrypt the files in local storage.

Creating a local certificate group

Local certificate groups are used to facilitate the configuration of profiles that are associated with a virtual server.

Before you begin, you must:

- Have Read-Write permission for System settings.
- Have already added the certificates to the local certificate store and intermediate CAs to the intermediate certificate store, and created an intermediate CA group.

To create a local certificate group:

1. Go to System > Certificate > Manage Certificates.
The configuration page displays the Local Certificate Group tab.
2. Click **Add** to display the configuration editor.
3. Complete the configuration as described in [Table 119](#).
4. Click **Save** when done.

Table 119: Local certificate group configuration

Settings	Guidelines
Group Name	Configuration name. Valid characters are A-Z, a-z, 0-9, _, and -. No spaces. The maximum length is 35 characters. After you initially save the configuration, you cannot edit the name.
Group Member	
Local Certificate	Select the certificate to add to the group.
Intermediate CA group	Select the Intermediate CA group to add to the local group. (Optional)
Default	Check this check box only if you want to make this local certificate the default for the group. Note: Only one local certificate can be set as the default in a group. If one local certificate has already been set as the default, you must disable (uncheck) it in order to set another one as the default. By default, the first local certificate in the group becomes the default if no local certificate is set as the default,

Importing intermediate CAs

An intermediate CA store is for the intermediate CA certificates that the backend servers would ordinarily use to complete the chain of server certificates, if any. HTTPS transactions use intermediate CAs when the server certificate is signed by an intermediate certificate authority (CA) rather than a root CA.

Before you begin, you must:

- Have Read-Write permission for System settings.
- Know the URL of an SCEP server or have downloaded the certificate and key files and be able to browse to them so that you can upload them.

To import an intermediate CA:

1. Go to System > Certificate > Manage Certificates.
2. Click the **Intermediate CA** tab.
3. Click **Import** to display the configuration editor.
4. Complete the configuration as described in [Table 120](#).

5. Click **Save** when done.
6. Repeat Steps 3 through 5 to import as many intermediate CAs as needed.

Table 120: Intermediate CA import configuration

Settings	Guidelines
Certificate Name	Configuration name. Valid characters are A-Z, a-z, 0-9, _, and -. No spaces. The maximum length is 35 characters. After you initially save the configuration, you cannot edit the name.
Import Method	<ul style="list-style-type: none"> • SCEP—Use Simple Certificate Enrollment Protocol. SCEP allows routers and other intermediary network devices to obtain certificates. • File—Upload a file.
SCEP	
SCEP URL	Specify the URL of the SCEP Server.
CA Identifier	Enter the identifier of the CA on the SCEP server, if applicable.
File	
Certificate File	Browse for and upload the the certificate file on the local machine.
Key File	Browse for the corresponding PEM key file that you want to upload. Note: Both a certificate file and key file are required for the intermediate CA used in SSL decryption by the forward proxy.
Password	Password to encrypt the files in local storage.

Creating an intermediate CA group

You select an intermediate CA group configuration object in the local certificate group, so you should configure in the group all the Intermediate CAs that would be needed by the backend servers that belong to a single virtual server.

Before you begin:

- You must have Read-Write permission for System settings.
- You must have already added the Intermediate CAs to the Intermediate CA certificate store.

To create an Intermediate CA group:

1. Go to System > Certificate > Manage Certificates.
2. Click the **Intermediate CA Group** tab.
3. Click **Add** to display the configuration editor.
4. Complete the configuration as described in [Table 121](#).
5. Save the configuration.

Table 121: Intermediate CA group configuration

Settings	Guidelines
Group Name	Configuration name. Valid characters are A-Z, a-z, 0-9, _, and -. No spaces. The maximum length is 35 characters. After you initially save the configuration, you cannot edit the name.
Group Member	
Intermediate CA	Select the Intermediate CA to add to the group,
Default	Check this check box only if you want to make this intermediate CA the default for the group. Note: Only one intermediate CA can be set as the default in an intermediate CA group. If one intermediate CA has already been set as the default, you must disable (uncheck) it in order to set another one as the default. By default, the first intermediate CA in a group becomes the default if no intermediate CA is set as the default,

Validating certificates

This section discusses the ways to validate client certificates and real server certificates from within the FortiADC system. It covers the following topics:

- [Importing CAs](#)
- [Creating a CA group](#)
- [Importing remote certificates](#)
- [Importing CRLs](#)
- [Adding OCSPs](#)
- [Validating certificates](#)

Configure a certificate verification object

To be valid, a client certificate must meet the following criteria:

- Must not be expired or not yet valid
- Must not be revoked by either certificate revocation list (CRL) or, if enabled, online certificate status protocol (OCSP)
- Must be signed by a certificate authority (CA) whose certificate you have imported into the FortiADC appliance

Certificate verification rules specify the CA certificates to use when validating client certificates, and they specify a CRL and/or OCSP server, if any, to use for certificate revocation checking.

You select a certificate verification configuration object in the profile configuration for a virtual server or in a real-server-SSL profile. If the client presents an invalid certificate during the authentication phase of a SSL/TLS session initiation, the FortiADC system will not allow the connection.

Before you begin:

- You must have Read-Write permission for System settings.
- You must have already created CA, OCSP or CRL configuration.

After you have configured a certificate verification object, you can include it in a virtual server profile or a Real Server SSL Profile, and it will be used to validate certificates presented to FortiADC.

To configure a certificate verification object:

1. Go to System > Certificate > Verify.
2. Click **Add** to display the configuration editor.
3. Complete the configuration as described in [Table 122](#).
4. Click **Save** when done. The newly certificate verification object appears on the Verify page.
5. Click the Edit icon in the far-right column (or double-click the entry) to open the configuration editor.
6. In the Group Member panel, select the CA, OCSP, or CRL of interest.
7. Click **Save** when done.

Table 122: Certificate verify configuration

Settings	Guidelines
Name	Enter a unique name for the certificate verification object that you are creating. Valid characters are A-Z, a-z, 0-9, <code>_</code> , and <code>-</code> . The maximum length is 35 characters. No space is allowed.
<code>verify-depth</code>	Note: CLI only. The default value is 1, but you may select any value from 0 to 255.
<code>customize-error-ignore</code>	Note: This option is available from the CLI only. Enable or disable <code>customize-error-ignore</code> . The option is disabled by default. If it's enabled, you are required to select the <code>ca-ignore-errors</code> and <code>cert-ignore-errors</code> , as described below.

Settings	Guidelines
<code>ca-ignore-errors</code>	<p>Note: CLI only. When <code>customize-error-ignore</code> is enabled, the following options become available for you to choose from:</p> <ul style="list-style-type: none">• <code>UNABLE_TO_GET_ISSUER_CERT</code>• <code>UNABLE_TO_GET_CRL</code>• <code>CERT_NOT_YET_VALID</code>• <code>CERT_HAS_EXPIRED</code>• <code>CRL_NOT_YET_VALID</code>• <code>CRL_HAS_EXPIRED</code>• <code>DEPTH_ZERO_SELF_SIGNED_CERT</code>• <code>SELF_SIGNED_CERT_IN_CHAIN</code>• <code>UNABLE_TO_GET_ISSUER_CERT_LOCALLY</code>• <code>UNABLE_TO_VERIFY_LEAF_SIGNATURE</code>• <code>CERT_CHAIN_TOO_LONG</code>• <code>INVALID_CA</code>• <code>INVALID_PURPOSE</code>• <code>CERT_UNTRUSTED</code>• <code>CERT_REJECTED</code> <p>Note: If <code>customize-error-ignore</code> is disabled (by default), the CLI shows the following:</p> <pre>ca-ignore-errors: UNABLE_TO_GET_ISSUER_CERT UNABLE_TO_GET_CRL CERT_UNTRUSTED</pre>

Settings	Guidelines
cert-ignore-errors	<p>Note: CLI only. When <code>customize-error-ignore</code> is enabled, the following options become available for you to choose from:</p> <ul style="list-style-type: none"> • <code>UNABLE_TO_GET_ISSUER_CERT</code> • <code>UNABLE_TO_GET_CRL</code> • <code>CERT_NOT_YET_VALID</code> • <code>CERT_HAS_EXPIRED</code> • <code>CRL_NOT_YET_VALID</code> • <code>CRL_HAS_EXPIRED</code> • <code>DEPTH_ZERO_SELF_SIGNED_CERT</code> • <code>SELF_SIGNED_CERT_IN_CHAIN</code> • <code>UNABLE_TO_GET_ISSUER_CERT_LOCALLY</code> • <code>UNABLE_TO_VERIFY_LEAF_SIGNATURE</code> • <code>CERT_CHAIN_TOO_LONG</code> • <code>INVALID_CA</code> • <code>INVALID_PURPOSE</code> • <code>CERT_UNTRUSTED</code> • <code>CERT_REJECTED</code> <p>Note: If <code>customize-error-ignore</code> is disabled (by default), the CLI shows the following:</p> <pre>cert-ignore-errors: UNABLE_TO_GET_CRL</pre>
Group Member	
CA	Select a CA (Required).
OCSP	Select an OCSP (Optional).
CRL	Select a CRL (Optional).

Importing CRLs

A certificate revocation list (CRL) is a file that contains a list of revoked certificates with their serial numbers and their revocation dates. The file also contains the name of the issuer of the CRL, the effective date, and the next update date. By default, the shortest validity period of a CRL is one hour.

Some potential reasons for certificates to be revoked include:

- A CA server was hacked and its certificates are no longer trustworthy.
- A single certificate was compromised and is no longer trustworthy.
- A certificates has expired and is not supposed to be used past its lifetime.

You can either upload a CRL file from your local machine or specify the URL of the CRL file



Online Certificate Status Protocol (OCSP) is an alternative to CRL. OCSP is useful when you do not want to deploy CRL files, for example, or want to avoid the public exposure of your PKI structure. For more information, see [Adding OCSPs](#).

Before you begin, you must:

- Have Read-Write permission for System settings.
- Know the URL of a CRL server or have the CRL files downloaded onto your local machine.

To import a CRL file:

1. Go to System > Certificate > Verify.
2. Click the **CRL** tab.
3. Click **Import** to display the configuration editor.
4. Complete the configuration as described in [Table 123](#).
5. Click **Save** when done.
6. Repeat Steps 3 through 5 to import as many CRLs as needed.

Table 123: CRL configuration

Settings	Guidelines
Name	Configuration name. Valid characters are A-Z, a-z, 0-9, _, and -. No spaces. The maximum length is 35 characters. After you initially save the configuration, you cannot edit the name.
Import Method	
HTTP	If selected, FortiADC will download the CRL file from an HTTP server. You must specify the HTTP URL.
SCEP	If selected, FortiADC will download the CRL file from an SCEP server. You must specify the SCEP URL.
File	If selected, you will need to browse for the CRL file on your local machine and upload it into FortiADC.

Adding OCSPs

FortiADC supports the validation of client digital certificates using Online Certificate Status Protocol (OCSP). In such a configuration, FortiADC contacts the OCSP Responder (i.e., the certificate management system), which maintains the current revocation status information of client certificates or backend server certificates, to determine the current status of digital certificate presented to it. It can then decide whether to allow or block the TLS/SSL connections, based on the status of the client certificates provided by the OCSP Responder.

OCSP enables you to validate certificate status by real-time online query, rather than by importing certificate revocation list (CRL) files. Since distributing and installing CRL files can be a considerable burden in large organizations, and because delay between the release and install of the CRL represents a vulnerability window, this can often be preferable.

During the process of TLS/SSL handshake, FortiADC will send an OCSP status request for the client certificate or backend server certificate to the OCSP Responder. The OCSP Responder then verifies whether the status request contains the information required to identify the certificate and returns a signed response with the status of the inquired certificate, which could be one of the following:

- Good = The certificate has not yet been revoked.
- Revoked = The certificate has been revoked.
- Unknown = The OCSP Responder has no information about the requested certificate, and therefore is able to determine its status.

Note: FortiADC only accepts client certificates in "Good" status as determined by the OCSP Responder as valid.

To use OCSP queries, you must first install the certificates of trusted OCSP servers.

Before you begin, you must:

- Have Read-Write permission for System settings.
- Know the URL of an OCSP server
- Have downloaded the certificate and key files and be able to browse to them so that you can upload them.
- Have already imported the OCSP signing certificates into FortiADC. See [Importing remote certificates](#) and [Creating a CA group](#).

To add an OCSP:

1. Go to System > Certificate > Verify.
2. Click the **OCSP** tab.
3. Click **Add** to display the configuration editor.
4. Complete the configuration as described in [Table 124](#).
5. Click **Save** when done.
6. Repeat Steps 3 through 5 to add as many OCSPs as needed.

Table 124: OCSP certificate configuration

Settings	Guidelines
Name	Enter a unique name for the client certificate validation object that uses OCSP. Valid characters are A-Z, a-z, 0-9, _, and -. The maximum length is 35 characters. No space is allowed.
OCSP URL	Specify the URL of the OCSP Responder.

Settings	Guidelines
Verify Others	<p>Upon receiving the OCSP response from the OCSP server, FortiADC first performs OCSP basic verify to validate the OCSP responder's signature.</p> <p>Enable (default)—When Verify Others is enabled, you must select a remote certificate (see Remote Certificates below). The OCSP basic verify succeeds when the selected remote certificate matches the OCSP response signature, otherwise, the OCSP basic verify will fail and the TLS/SSL connection will be terminated.</p> <p>Disable—When Verify Others is disabled, you must select a CA Group. the OCSP basic verify will be carried out in the following sequence:</p> <ol style="list-style-type: none"> 1. The OCSP response signing certificate must be one of the certificates in the CA group or a certificate issued by one of the certificates in the CA group. Also, the certificates must form a chain from the OCSP signing certificate all the way to a self-signed root CA. Otherwise, the OCSP basic verify will fail. 2. If Step 1 (above) is successful, the validation will proceed like this: If the Criteria Check field is selected (enabled), then the OCSP signing certificate can be either the issuing CA of the certificate whose status FortiADC must validate, or a dedicated OCSP signing certificate issued by this issuing CA. The validation succeeds if this criterion is met. Otherwise, the validation process will move onto Step 3 (below). 3. If the OCSP signing certificate is issued by one of the certificates in the CA group, but is not a dedicated OCSP signing certificate, then the validation will proceed like this: If the root CA of this OCSP signing certificate is a trusted self-signed root CA and the "Accept Trusted Root CA" field is selected (enabled), then the validation will succeed. Otherwise, the validation will fail.
Remote Certificates	<p>Select the client certificate of which you'd like to verify the signature of the OCSP Responder that signs it. Note: This option is applicable only when Verify Others is enabled. You MUST select a remote certificate which must have been imported into FortiADC ahead of time. See Importing remote certificates.</p>
CA Chain	<p>Click the down arrow and select a CA group from the list menu. Note: This becomes available only when Verify Others is disabled. In that case, you must select a CA chain (i.e., CA group). It's highly recommended that you have CA groups configured in advance to use this option. See Creating a CA group.</p>
Criteria Check	<p>Enable/Disable issuer-criteria check. Note: This option comes in hand in hand with CA Chain, and is only available when Verify Others is disabled (see Verify Others above). It is enabled by default, but you can uncheck it if you do not want to validate the certificate issuer's identity.</p>
Accept Trusted Root CA	<p>Enable/Disable accept trusted root CA. Note: This option becomes available only when Criteria Check is enabled (see Criteria Check above). It is enabled by default, in which case FortiADC will accept trusted root CA in the validation process. Uncheck it if you do not want to use this feature.</p>

Importing OCSP signing certificates

OCSP signing certificates are certificates with no private keys. For dynamic certification revocation, you must verify them through an OCSP server. This option allows you to import remote (OCSP) certificates into FortiADC

and use them to verify the OCSP response signature.

Before you begin, you must:

- Have Read-Write permission for System settings.
- Have the remote certificates downloaded onto your local machine so that you can upload it to FortiADC.

To import an OCSP-signing certificate:

1. Go to System > Certificate > verify.
2. Click the **OCSP Signing Certificate** tab.
3. Click **Import** to display the configuration editor.
4. Complete the configuration as described in [Table 125](#).
5. Click **Save** when done.
6. Repeat Steps 3 through 5 to import as many remote certificates as needed.

Table 125: Importing an OCSP signing certificate

Settings	Guidelines
Name	Enter a unique name for the remote certificate you want to import. Valid characters are A-Z, a-z, 0-9, _, and -. The maximum length is 35 characters. No space is allowed.
OCSP Signing Certificates	Browse for and upload the remote certificate file of interest.

Once an OCSP signing certificate has been uploaded into FortiADC, the name of the certificate file shows up under the Remote tab. You can view or remove the certificate from this page using the corresponding icons in the far-right column of the page.

Importing CAs

The certificate authority (CA) store is used to authenticate the certificates of other devices. When the FortiADC system is presented with a certificate, it examines the CA's signature, comparing it with the copy of the CA's certificate already imported into the CA store. If the public key matches the private key, the client's or device's certificate is considered legitimate.

In web browsers, the CA store includes trusted root CAs that can be used to establish trust with servers that have certificates signed by the issuing CAs. In an SSL forward proxy deployment, FortiADC acts as a proxy for the client, so you might want to import client browser CAs, create a CA group, and create a certificate verification policy to verify server certificates against this group. You can examine the CA store in common web browsers to come up with a good list of CAs to download and then import. The following list has links for some common web browsers:

- Apple iOS: <https://support.apple.com/en-us/HT204132>
- Google Chrome and Mozilla Firefox: <https://wiki.mozilla.org/CA:IncludedCAs>
- Microsoft Internet Explorer: <https://technet.microsoft.com/en-us/library/dn265983.aspx>

You must do one of the following:

- Import the certificates of the signing CA and all intermediate CAs to FortiADC's store of CA certificates.
- In *all* personal certificates, include the full signing chain up to a CA that FortiADC knows in order to prove that the clients' certificates should be trusted.

- If the signing CA is not known, that CA's own certificate must likewise be signed by one or more other intermediary CAs, until both the FortiADC appliance and the client or device can demonstrate a signing chain that ultimately leads to a mutually trusted (shared "root") CA that they have in common. Like a direct signature by a known CA, this proves that the certificate can be trusted.

Before you begin, you must:

- Have Read-Write permission for System settings.
- Know the URL of an SCEP server or have downloaded the certificate and key files and be able to browse to them so that you can upload them.

To import a CA:

1. Go to System > Certificate > Verify.
2. Click the **CA** tab.
3. Click **Import** to display the configuration editor.
4. Complete the configuration as described in [Table 126](#).
5. Click **Save** when done.
6. Repeat Steps 3 through 5 to import as many CAs as needed.

Table 126: CA import configuration

Settings	Guidelines
Certificate Name	Configuration name. Valid characters are A-Z, a-z, 0-9, _, and -. The maximum length is 35 characters. No space is allowed.
Import Method	<ul style="list-style-type: none"> • SCEP—Use Simple Certificate Enrollment Protocol. SCEP allows routers and other intermediary network devices to obtain certificates. • File—Upload a file.
SCEP	
SCEP URL	Enter the URL of the SCEP server.
CA Identifier	Enter the identifier for a specific CA on the SCEP server.
File	
Local PC	Browse for the certificate file on the local machine and upload it to FortiADC.

Creating a CA group

CA groups are only used to verify the signature of the OCSP Responder.

Include in the CA group all of the CAs for the pool of backend servers to be associated with a single virtual server.

Before you begin, you must:

- Have Read-Write permission for System settings.
- Have already added the CAs to the CA certificate store.

To create a CA group:

1. Go to System > Certificate > Verify.
2. Click the **CA Group** tab.
3. Click **Add** to display the configuration editor.
4. Name the CA group and click **Save** when done. The new CA group appears on the CA Group page.
5. Click the Edit icon in the far-right column (or double-click the CA group) to bring up the configuration editor.
6. Click **Add**.
7. Complete the configuration as described in [Table 127](#).
8. Click **Save** when done.
9. Repeat Steps 6 through 8 to add as many CAs to the group as needed.

Table 127: CA group configuration

Settings	Guidelines
Group Name	Specify a unique name for the CA group that you are creating. Valid characters are A-Z, a-z, 0-9, _, and -. The maximum length is 35 characters. No space is allowed.
Group Member	
CA	Click the down arrow and select the desired CA from the list menu to add to the group.

HSM Integration

A hardware security module (HSM) is a dedicated device for managing digital keys and performing cryptographic operations. It can be a plug-in card or an external device directly connected to a computer or network server. Purposefully designed to protect the crypto-key life cycle, HSMs have been used by some of the world's most security-conscious entities to protect their cryptographic infrastructure by securely managing, processing, and storing cryptographic keys inside a hardened, tamper-resistant device.

Enterprises use HSMs to safeguard their online transactions, identities, and applications because of their strengths in securing cryptographic keys and provisioning encryption, decryption, authentication, and digital signing services for a wide range of applications.

Integrating FortiADC with SafeNet Network HSM

Starting from Version 4.7.2, FortiADC has integrated with SafeNet Network HSM. It enables you to retrieve a per-connection, SSL session key from the HSM server instead of loading the private key and certificate stored on FortiADC.

Integration of FortiADC with SafeNet Network HSM requires specific configuration steps on both appliances:

On the HSM appliance:

- Create one or more HSM partitions for FortiADC
- Send the FortiADC client certificate to the HSM
- Register the FortiADC HSM client to the partition(s)
- Retrieve the HSM server certificate

On the FortiADC appliance:

- Configure communication with the HSM server, including using the server and client certificates to register FortiADC as a client of the HSM
- Generate a certificate-signing request (CSR) that includes the HSM's configuration information
- Upload the signed certificate to FortiADC

It must be noted that

- HSM support is disabled on FortiADC by default. You must enable it via the CLI for the feature to become available on FortiADC GUI. Use the following CLI commands to enable HSM support:

```
config system global
set hsm enable
```

- You must also have the HSM server certificate available on your local PC or a network drive.
- With the 4.7.2 release, HSM integration only supports active-passive HA configuration.
- In Version 4.7.2, HSM partition is a global configuration that can be used from individual VDOMs.

The following instructions assume that you have (1) HSM support enabled on FortiADC and (2) access to the HSM server certificate from your PC.

Preparing the HSM appliance

FortiADC 4.7.2 only supports SafeNet Network HSM. Before starting to configure the FortiADC-HSM integration, you must configure the SafeNet Network HSM first using the following steps:

1. On the SafeNet Network HSM, use the `partition create` command to create and initialize a new HSM partition that uses password authentication.
Note: This is the partition FortiADC uses on the HSM server. You can create more than one partition, but all the partitions are assigned to the same client. For more information, see the HSM-related documentation.
2. Use the SCP utility and the following command to send the FortiADC client certificate to the HSM:
`scp <fortiadc_ip>.pem admin@<hsm_ip>:`
3. Using SSH, connect to the HSM server using the admin account. Then, use the following command to register a client for FortiADC on the HSM server:
`lunash:> client register -c <client_name> -ip <fortiadc_ip>, where <client_name> is the name you specify that identifies the client.`
4. Use the following command to assign the client you registered to the partition you created earlier:
`lunash:> client assignPartition -client <client_name> -partition <partition_name>`
You can verify the assignment using the following command:
`lunash:> client show -client <client_name>`
5. Repeat the client assignment process for any additional partitions you created for FortiADC.
6. Use the SCP utility and the following command to retrieve the server certificate file from the HSM server:
`scp <hsm_username>@<hsm_ip>:server.pem /usr/lunasa/bin/server_<hsm_ip>.pem`
7. On the FortiADC GUI, navigate to **System>HSM** to bring up the HSM configuration page.
8. Complete the HSM configuration as described in the table below. Then move on to "[HSM Integration](#)" on page 333.

Table 128: HSM Configuration Parameters

Parameter	Description
Configuration	Complete the following entries or selections to configure the FortiADC-HSM integration.

Parameter	Description
IP	Enter the IP address of the HSM server.
Port	Specify the port via which FortiADC establishes an NTLS connection with the HSM server. The default value is 1792.
Timeout	Specify a timeout value for the connection between FortiADC and the HSM server. The default is 20000. Valid values range from 5000 to 20000 milliseconds.
Interface	Select an outgoing interface (port) on FortiADC. Note: FortiADC uses this interface (port) to generate a local certificate based on the IP address of the interface (e.g., port1).
Upload Server Certificate File	Click Browse to browse for the server certificate file that you retrieved earlier.
Download Client Certificate File	Click Download to retrieve the client certificate file that you sent to the HSM server earlier to make it available for the registration process. Note: The certificate's common name is the main IP address of the outgoing interface. The client and sever certificates are used in SSL connection between FolritADC and the HSM server.
Register	Click this button to register FortiADC as a client of the HSM sever using the specified server and client certificates. Note: This action generates a config file, e.g., <code>/example.conf</code>
Unregister	Click this button to clear all HSM-related configurations on the back-end.
Partition	Click Add to create partition or Delete to remove a selected partition. Note: FortiADC 4.7.2 can accept only one partition. Once a partition is added, the Register and Unregister buttons become dimmed out, meaning you cannot make any change to the HSM configuration. To eidt the HSM configuration, you must delete the partition first.
Partition Name	Specify the name of a partition to which the FortiADC HSM client is assigned.
Password	Specify the password for the partition.

Note: When configure your CSR to work with an HSM, the CSR generation process creates a private key on both the HSM and the FortiADC. The private key on the HSM is the "real" key that secures communication when FortiADC uses the signed certificate. The key found on the FortiADC is used when you upload the certificate to FortiADC.

Generating a certificate-signing request on FortiADC

Once you have completed configuring the HSM server, you must generate a certificate-signing request which references the HSM connection and partition from inside FortiADC.

To generate a certificate-signing request:

1. On the FortiADC GUI, navigate to **System > Manage Certificates > Local Certificate**.
2. Click **Generate** to bring up the Local Certificate configuration page.
3. Configure the certificate-signing request as described in the table below. Then move on to ["HSM Integration" on page 335](#).

Table 129: Generating a certificate-signing request

Parameter	Description
Generate Certificate Signing Request	Complete the following entries or selections to configure the FortiADC-HSM integration.
Certificate Name	Specify a name for the certificate request, e.g., www.example.com. This can be the name of your web site.
Subject Information	Specify the information that the certificate is required to contain in order to uniquely identify the FortiADC appliance. This area varies depending on the ID Type you select.
ID Type	<p>Select the type of identifier to use in the certificate to identify the FortiADC appliance:</p> <ul style="list-style-type: none"> • Host IP — Select this option if the FortiADC appliance has a static IP address, and then enter the public IP address of the FortiADC appliance in the IP field. If the FortiADC appliance does not have a public IP address, use Domain Name or Email instead. See below. • Domain Name — Select this option if the FortiADC appliance has a static IP address and subscribes to a dynamic DNS service. Enter the FQDN of the FortiADC appliance, such as www.example.com, in the Domain Name field, but do NOT include the protocol specification (http://) or any port number or path names. • Email — Select this option if the FortiADC appliance does not require either a static IP address or a domain name. Enter the email address of the owner of the FortiADC appliance in the Email field. <p>The ID type you can select varies by whether or not your FortiADC appliance has a static IP address, a fully-qualified domain name (FQDN), and by the primarily intended use of the certificate. For example, if your FortiADC appliance has both a static IP address and a domain name, but you will primarily use the local certificate for HTTPS connections to the web UI by the domain name of the FortiADC appliance, you might prefer to generate a certificate based upon the domain name of the FortiADC appliance rather than its IP address. Depending on your choice for ID Type, the other options may vary.</p>

Parameter	Description
IP	<p>Note: This option appears only if the ID Type is Host IP.</p> <p>Enter the static IP address of the FortiADC appliance, such as 10.0.0.1. The IP address must be the one visible to clients. Usually, this should be its public IP address on the Internet, or a virtual IP that you use NAT to map to the appliance's IP address on your private network.</p>
Domain Name	<p>Note: This option appears only if the ID Type is Domain Name.</p> <p>Enter the fully qualified domain name (FQDN) of the FortiADC appliance, such as www.example.com. The domain name must resolve to the static IP address of the FortiADC appliance or a protected server.</p>
Email	<p>Note: This option appears only if the ID Type is Email.</p> <p>Enter the email address of the owner/user of the FortiADC appliance, such as admin@example.com.</p>
Distinguished Information	The following information is OPTIONAL in the certificate; it is NOT required.
Organization unit	<p>Enter the name of your organizational unit (OU), such as the name of your department.</p> <p>To enter more than one OU name, click the + icon, and enter each OU in each separate field.</p>
Organization	Enter the legal name of your organization.
Locality(City)	Enter the name of the city or town where the FortiADC appliance is deployed.
State/Province	Enter the name of the state or province where the FortiADC appliance is deployed.
Country/Region	Select the name of the country where the FortiADC appliance is deployed.
Email	Enter an email address that may be used for contact purposes, such as admin@example.com.
Key Information	Enter the information pertinent to the key.
Key Type	<p>This field shows the type of algorithm used to generate the key.</p> <p>Note: It's read-only and cannot be changed. FortiADC 4.7.2 supports RSA key type only.</p>

Parameter	Description
Key Size	<p>Select one of the following key sizes:</p> <ul style="list-style-type: none"> • 512 bit • 1024 bit • 1536 bit • 2048 bit • 4096 bit <p>Note: Larger keys may take longer to generate, but provide better security.</p>
HSM	<p>Select this option if the private key for the connections is provided by an HSM appliance instead of FortiADC.</p> <p>Note: This option is available only if you have enabled HSM via the CLI using the <code>config system global</code> command. For more information, see "HSM Integration" on page 330.</p>
Partition Name	<p>Enter the name of the partition where the private key for this certificate is located on the HSM server.</p> <p>Note: This option becomes available only when HSM is selected. See above.</p>
Enrollment Information	
Enrollment Method	<p>Select either of the following:</p> <ul style="list-style-type: none"> • File Based —If selected, you must manually download and submit the resulting certificate signing request (.csr) file to a certificate authority (CA) for signing. Once signed, you need to upload the local certificate. <i>This is the only enrollment method if HSM is selected.</i> • Online SCEP — If selected, the FortiADC appliance will automatically use HTTP to submit the certificate-signing request to the simple certificate enrollment protocol (SCEP) server of a CA, which will validate and sign the certificate. <p>Note: For this selection, two more options appear: CA Server URL and Challenge Password. <i>This option is not available if HSM is selected.</i></p>

Downloading and uploading the certificate request (.csr) file

Normally, when generating a certificate-signing request, the FortiADC appliance creates a private and public key pair. The generated request includes the public key of the FortiADC appliance and information such as the FortiADC appliance's IP address, domain name, or email address. The FortiADC appliance's private key remains confidential on the FortiADC appliance. The Status column of the entry is PENDING.

If you configured your CSR to work with the FortiADC-HSM integration, the CSR generation process creates a private key both on the HSM and on FortiADC appliances. The private key on the HSM is used to secure communication when FortiADC uses the certificate. The FortiADC private key is used when you upload the certificate to FortiADC.

After you have submitted a certificate-signing request from inside FortiADC as discussed above, you must go back to the **System > Management Certificates > Local Certificate** page to download the certificate request (.csr) file, and then upload that file to your certificate authority (CA) by taking the following steps:

1. On the **System > Manage Certificates > Local Certificate** page, locate the entry of the certificate request.
2. Click the **Download** icon.
Note: The time it takes to download the certificate request (.csr) file varies, depending on the size of the file and the speed of your network connection. After the file is downloaded, save it at a location on your machine.
3. Upload the certificate request (.csr) file to your CA.
Note: Upon receiving the certificate request file, the CA will verify the information in the certificate, give it a serial number and an expiration date, and sign it with the public key of the CA.
4. If you are not using a commercial CA whose root certificate is already installed by default on web browsers, download your CA's root certificate, and then install it on all computers that will be connecting to your FortiADC appliance.
Note: You must have the certificate installed on the computers. Otherwise, they may not trust your new certificate. After you have received the signed certificate from the CA, upload it to FortiADC, as discussed below.

Uploading the server certificate to FortiADC

You must have the Read and Write permission to upload server certificates to the FortiADC appliance.

To upload the server certificate to FortiADC:

1. On the FortiADC GUI, navigate to the **System > Manage Certificates > Local Certificate** page.
2. Click **Import**.
3. Make the selections as described in the table below, and click **Save**.

Table 130: Uploading a server certificate-

Parameter	Description
Type	<p>Click the down arrow and select one of the following options from the drop-down menu:</p> <ul style="list-style-type: none"> • Local Certificate—Use this option only if you have a CA-signed certificate that was originated from a CSR generated in FortiADC . See HSM Integration on page 330. Note: It is important to make sure that the load-balancer (FortiADC appliance) you use to import a local certificate is the same appliance where the CSR was generated because it is where the key matching the certificate resides. The import operation will fail without the matching key on the same hardware system. • PKCS12 Certificate—Use this option only if you have a PKCS #12 password-encrypted certificate with its key in the same file. • Certificate—Use this option only if you have a certificate and its key in separate files. <p>Note: Additional fields are displayed depending on your selection.</p>
Certificate File	Click Browse to locate the certificate file that you want to upload.

Parameter	Description
Certificate Name	The name of the certificate. Note: This field applies when Type is Certificate or PKCS12 .
Key File	Click Browse to locate the key file that you want to upload with the certificate. Note: This option is available only if Type is Certificate .
Password	Enter the password used to encrypt the server certificate file. Note: This enables FortiADC to decrypt and install the certificate. This option is available only if Type is Certificate or PKCS12 Certificate .

Once a certificate is uploaded to FortiADC, you can use in a policy or server pool configuration. For more information, see

Chapter 13: Logging and Reporting

This chapter includes the following topics:

- Using the event log
- Using the aggregate log
- Using the traffic log
- Configuring local log settings
- Configuring syslog settings
- Enabling real-time statistics
- Configuring high speed logging
- Configuring alert email settings
- Configuring reports
- Configuring Report Queries
- Configuring fast reports
- Viewing reports

Using the event log

The Event Log table displays logs related to system-wide status and administrator activity.

Figure 56 shows the Event log table. By default, the log is filtered to display configuration changes, and the table lists the most recent records first.

You can use the following category filters to review logs of interest:

- Configuration—Configuration changes.
- System—System operations, warnings, and errors.
- Admin—Administrator actions.
- User—Authentication results logs.
- Health Check—Health check results and client certificate validation check results.
- SLB—Notifications, such as connection limit reached.
- LLB—Notifications, such as bandwidth thresholds reached.
- GLB—Notifications, such as the status of associated local SLB and virtual servers.
- Firewall—Notifications, such as SNAT source IP pool is using all of its addresses.

Figure 56: Event log

Date	Time	Log Level	User	Action	Message
2017-05-16	11:32:18	Information	daemon_admin	edit	changed settings 'logonly-response' (0->1494959536) for 'system scripting' on domain 'root';
2017-05-16	09:51:38	Information	admin	edit	changed settings 'port1' for 'system interface' on domain 'root';
2017-05-11	00:22:50	Information	admin	del	deleted an entry 'nat_scv' for 'firewall nat-snat' on domain 'root';

Within each category, you can use Filter Setting controls to filter the table based on the values of matching data.

You can use the **Download** link to download the logs. Filters are applied to the set that is collected for download.

Table 131: Filter settings

Category Filters	Data Filters
Configuration	<ul style="list-style-type: none"> • Date • Time • Priority (Log Level) • User • Action
System	<ul style="list-style-type: none"> • Date • Time • Priority (Log Level) • Submod • User • Action • Status
Admin	<ul style="list-style-type: none"> • Date • Time • Priority (Log Level) • User • Action • Status
User	<ul style="list-style-type: none"> • Date • Time • Log Level • User • Action • Status
Health Check	<ul style="list-style-type: none"> • Date • Time • Priority (Log Level) • Module • Policy • Group • Member • Status

Category Filters	Data Filters
SLB, LLB, GLB, Fire-wall	<ul style="list-style-type: none"> • Date • Time • Priority (Log Level) • Module • Policy • Group • Member • Status • Action

The last column in each table includes a link to log details.

Before you begin:

- You must have Read-Write permission for Log & Report settings.

To view and filter the log:

1. Go to Log & Report > Log Browsing.
The log page displays the Event Logs tab.
2. Select the category of interest.
3. Click **Filter Setting** to display the filter tools.
4. Use the tools to filter on key columns and values.
5. Click **OK** to apply the filter and redisplay the log.

Table 132 to Table 137 list the log columns for the event log types in the order in which they appear in the log.

Table 132: Event log — Config

Column	Example	Description
date	date=2014-12-01	Log date.
time	time=15:50:37	Log time.
log_id	log_id=0000000085	Log ID.
type	type=event	Log type.
subtype	subtype=config	Log subtype.
pri	pri=information	Log level.
vd	vd=root	Virtual domain.
msg_id	msg_id=522000	Message ID.
user	user=admin	User that performed the operation.

Column	Example	Description
ui	ui=GUI(172.30.144.8)	User interface from which the operation was performed.
action	action=add	Administrator action: add, edit, delete.
cfgpath	cfgpath=firewall qos-queue	Configuration that was changed.
cfgobj	cfgobj=name	Configuration setting changed.
cfgattr	cfgattr=queue	Configuration value changed.
logdesc	logdesc=Change the configuration	A column added for compatibility with FortiAnalyzer.
msg	msg=added a new entry 'queue' for "firewall qos-queue" on domain "root"	Log message.

Table 133: Event log — System

Column	Example	Description
date	date=2014-12-01	Log date.
time	time=16:00:09	Log time.
log_id	log_id=0003000011	Log ID.
type	type=event	Log type.
subtype	subtype=system	Log subtype.
pri	pri=error	Log level.
vd	vd=root	Virtual domain.
msg_id	msg_id=522008	Message ID.
submod	submod=update	System submodule.
user	user=none	None.
ui	ui=none	None.
action	action=update	System action, such as (firmware) update, HA join and leave, and the like.
status	status=failure	Status message: success or failure.
logdesc	logdesc=Update FortiGuard	A column added for compatibility with FortiAnalyzer.
msg	msg=Update firmware	Log message (if any).

Table 134: Event log — Admin

Column	Example	Description
date	date=2014-12-01	Log date.
time	time=15:44:38	Log time.
log_id	log_id=0001016834	Log ID.
type	type=event	Log type.
subtype	subtype=admin	Log subtype.
pri	pri=information	Log level.
vd	vd=root	Virtual domain.
msg_id	msg_id=521996	Message ID.
user	user=admin	User that performed the operation.
ui	ui=GUI(172.30.144.8)	User interface from which the operation was performed.
action	action=logout	System action.
status	statue=success	Status message: success or failure.
reason	reason=none	Reason string (if any).
logdesc	logdesc=Admin login	A column added for compatibility with FortiAnalyzer.
msg	msg=User admin logout from GUI(172.30.144.8).	Log message.

Table 135: Event log — User

Column	Example	Description
date	date=2014-12-01	Log date.
time	time=15:44:38	Log time.
log_id	log_id=0001016834	Log ID.
type	type=event	Log type.
subtype	subtype=user	Log subtype.
pri	pri=information	Log level.

Column	Example	Description
vd	vd=root	Virtual domain.
msg_id	msg_id=521996	Message ID.
user	user=user1	User name.
usergrp	usergrp=customerABC	User group.
policy	policy=membersOnly	Authentication policy.
action	action=login	System action.
status	statue=success	Status message: success or failure.
reason	reason=none	Reason string (if any).
logdesc	logdesc=user auth	A column added for compatibility with FortiAnalyzer.
msg	msg=User admin logout from GUI(172.30.144.8).	Log message.

Table 136: Event log — Health Check

Column	Example	Description
date	date=2015-12-30	Log date.
time	time=12:07:47	Log time.
log_id	log_id=2002502	Log ID.
type	type=event	Log type.
subtype	subtype=health	Log subtype.
pri	pri=alert	Log level.
vd	vd=root	Virtual domain.
msg_id	msg_id=35661161	Message ID.
module	smodule=slb	System module: slb, llb.
policy	policy=HTTPS_VIP	Virtual server configuration to which the event applies.
group	group=test2	Real server pool group or link group.
member	member=1	Real server member ID or gateway ID.

Column	Example	Description
attrtype	attrtype=none	Attribute type (if any).
attrname	attrname=none	Attribute type (if any).
action	action=health_check	Type of message: health check.
status	status=failure	Health check result: success or failure.
logdesc	logdesc=SLB Virtual server change state	A column added for compatibility with FortiAnalyzer.
msg	msg=Virtual server HTTPS_VIP, status is down	Log message.

Table 137: Event log — SLB, LLB, GLB, Firewall

Column	Example	Description
date	date=2016-01-13	Log date.
time	time=08:30:12	Log time.
log_id	log_id=0005001704	Log ID.
type	type=event	Log type.
subtype	subtype=slb	Log subtype: dns (glb), slb, llb, fw.
pri	pri=alert	Log level.
vd	vd=root	Virtual domain.
msg_id	msg_id=115208	Message ID.
policy	policy=L7vs_tcps	Policy to which the event applies—the virtual server configuration name, for example.
group	group=none	Real server pool group or link group.
member	member=none	Real server member ID or gateway ID.
attrtype	attrtype=none	Additional configuration attributes, if applicable.
attrname	attrname=none	Additional configuration values, if applicable.
action	action=ssl	Module that took action.
status	status=failure	Status of action.

Column	Example	Description
logdesc	logdesc=SLB SSL Handshake	A column added for compatibility with FortiAnalyzer.
msg	msg=Client 31.1.1.103 failed to establish SSL connection with VS 41.1.1.123	Log message.



The value "none" appears in logs when the value is irrelevant to the status or action. For example, a health check log for a virtual server shows "none" in the Group and Member columns even though its real server pool and members are known—these details are just not relevant. Likewise, a health check log for a real server pool member shows "none" in the Policy column even though its virtual server is known.

Using the security log

The Security Log table displays logs related to security features.

Figure 57 shows the security log table. By default, the log is filtered to display IP Reputation logs, and the table lists the most recent records first.

You can use the following category filters to review logs of interest:

- IP Reputation—Traffic logged by the IP Reputation feature
- DoS—Traffic logged by the SYN Flood feature
- WAF—Traffic logged by the WAF feature
- Geo—Traffic logged by the Geo IP block list feature

Figure 57: Security log

Event Log Security Log Traffic Log Script Log							
<input type="radio"/> IP Reputation <input checked="" type="radio"/> DoS <input type="radio"/> WAF <input type="radio"/> GEO							
Filter Setting				Download		Refresh	
Date	Time	Count	Severity	Source	Destination	Action	
2016-11-30	17:31:57	0	high	0.0.0.0	10.1.50.201	deny	
2016-11-30	17:31:35	10	high	0.0.0.0	10.1.50.201	deny	
2016-11-30	17:31:35	1	high	0.0.0.0	10.1.50.201	deny	
2016-11-30	17:31:18	10	high	0.0.0.0	10.1.50.201	deny	
2016-11-30	17:31:18	1	high	0.0.0.0	10.1.50.201	deny	
2016-11-30	17:31:07	10	high	0.0.0.0	10.1.50.201	deny	
2016-11-30	17:30:11	9	high	0.0.0.0	10.1.50.201	deny	
2016-11-30	17:30:00	10	high	0.0.0.0	10.1.50.201	deny	
2016-11-30	17:29:56	312	high	0.0.0.0	10.1.50.201	deny	
2016-11-30	17:29:33	322	high	0.0.0.0	10.1.50.201	deny	

Within each category, you can use Filter Setting controls to filter the table based on the values of matching data:

- Date
- Time
- Proto
- Service
- Src
- Src_port
- Dst
- Dst_port
- Vs Name
- Action

The last column in each table includes a link to log details.

Before you begin:

- You must have Read-Write permission for Log & Report settings.

To view and filter the log:

1. Go to Log & Report > Log Browsing.
2. Click the **Security Logs** tab to display the attack log.
3. Click **Filter Settings** to display the filter tools.
4. Use the tools to filter on key columns and values.
5. Click **OK** to apply the filter and redisplay the log.

[Table 138](#) to [Table 141](#) list the log columns in the order in which they appear in the log.

Table 138: IP Reputation log

Column	Example	Description
date	date=2014-12-02	Log date.
time	time=10:27:01	Log time.
log_id	log_id=0200004230	Log ID.
type	type=attack	Log type: attack.
subtype	subtype=ip_reputation	Log subtype: ip_reputation.
pri	pri=warning	Log level.
vd	vd=root	Virtual domain.
msg_id	msg_id=13065998	Message ID.
count	count=1	For IP reputation, count=1.
severity	severity=high	Rule severity.

Column	Example	Description
proto	proto=6	Protocol.
service	service=http	Service.
src	src=4.4.4.4	Source IP address.
src_port	src_port=49301	Source port.
dst	dst=2.2.2.2	Destination IP address.
dst_port	dst_port=80	Destination port.
policy	policy=vs1	Virtual server name.
action	action=deny	Policy action.
srccountry	srccountry=cn	Location of the source IP address.
dstcountry	dstcountry=us	Location of the destination IP address.
msg	msg=msg	Security rule name, category, subcategory, and description of the attack.

Table 139: DoS log

Column	Example	Description
date	date=2014-12-02	Log date.
time	time=10:27:01	Log time.
log_id	log_id=0200004230	Log ID.
type	type=attack	Log type: attack.
subtype	subtype=synflood	Log subtype: synflood.
pri	pri=warning	Log level.
vd	vd=root	Virtual domain.
msg_id	msg_id=13065998	Message ID.
count	count=1	For DoS, number of timeouts sent per destination.
severity	severity=high	Always “high” for DoS.
proto	proto=0	Protocol.

Column	Example	Description
service	service=http	Service.
src	src=173.177.99.94	Source IP address.
src_port	src_port=49301	Source port.
dst	dst=10.61.2.100	Destination IP address.
dst_port	dst_port=80	Destination port.
policy	policy=unknown	For DoS, policy=unknown.
action	action=deny	Policy action.
srccountry	srccountry=cn	Location of the source IP address.
dstcountry	dstcountry=us	Location of the destination IP address.
msg	msg=msg	Security rule name, category, subcategory, and description of the attack.

Table 140: WAF log

Column	Example	Description
date	date=2015-07-22	Log date.
time	time=10:27:01	Log time.
log_id	log_id=0202008074	Log ID.
type	type=attack	Log type: attack.
subtype	subtype=waf	Log subtype: waf.
pri	pri=alert	Log level.
vd	vd=root	Virtual domain.
msg_id	msg_id=1512	Message ID.
count	count=1	Rule match count.
severity	severity=low	Rule severity.
proto	proto=6	Protocol.
service	service=http	Service.

Column	Example	Description
src	src=1.1.1.1	Source IP address.
src_port	src_port=34352	Source port.
dst	dst=2.2.2.2	Destination IP address.
dst_port	dst_port=80	Destination port.
policy	policy=vs1	Virtual server name.
action	action=pass	Policy action.
sigid	sigid=1	Attack signature ID.
subcat	subcat=waf_subtype	WAF module: waf_web_attack_signature, waf_url_access, waf_http_protocol_cont and waf_sql_xss_injection_detect.
http_host	http_host-t=192.168.1.140:8080	HTTP Host header in HTTP request. Maximum length is 64. Longer URIs are truncated and appended with . . .
http_url	http_url=/bigdata	URI in HTTP request. Maximum length is 128. Longer URIs are truncated and appended with . . .
pkt_hdr	pkt_hdr=header	Contents of the packet header that matched the attack signature.
srccountry	srccountry=Australia	Location of the source IP address.
dstcountry	dstcountry=France	Location of the destination IP address.
msg	msg="Find Attack ID: 1010010001 NAME: "HTTP Method Violation" CATEGORY: "HTTP Protocol Constraint" SUB_CATEGORY: "Request Method Rule""	Security rule name, category, subcategory, and description of the attack.

Table 141: Geo IP log

Column	Example	Description
date	date=2014-12-02	Log date.
time	time=10:27:01	Log time.
log_id	log_id=0200004230	Log ID.
type	type=attack	Log type: attack.

Column	Example	Description
subtype	subtype=geo	Log subtype: geo.
pri	pri=warning	Log level.
vd	vd=root	Virtual domain.
msg_id	msg_id=13065998	Message ID.
count	count=1	Rule match count.
severity	severity=high	Rule severity.
proto	proto=0	Protocol.
service	service=http	Service.
src	src=173.177.99.94	Source IP address.
src_port	src_port=49301	Source port.
dst	dst=10.61.2.100	Destination IP address.
dst_port	dst_port=80	Destination port.
policy	policy=vs1	Virtual server name.
action	action=deny	Policy action.
srccountry	srccountry=cn	Location of the source IP address.
dstcountry	dstcountry=us	Location of the destination IP address.
msg	msg=msg	Security rule name, category, subcategory, and description of the attack.

Using the traffic log

The Traffic Log table displays logs related to traffic served by the FortiADC deployment.

shows the Traffic log table. By default, the log is filtered to display Server Load Balancing - Layer 4 traffic logs, and the table lists the most recent records first.

You can use the following category filters to review logs of interest:

- SLB Layer 4—Traffic served by Layer-4 virtual servers
- SLB HTTP—Traffic served by virtual servers with HTTP profiles
- SLB TCPS—Traffic served by virtual servers with TCPS profiles
- SLB RADIUS—Traffic served by virtual servers with RADIUS profiles
- GLB—Traffic served by global load balancing policies

- SLB SIP—Traffic served by virtual servers with SIP profiles
- SLB RDP—Traffic served by virtual servers with RDP profiles
- SLB DNS —Traffic served by virtual servers with DNS profiles
- SLB RTSP —Traffic served by virtual servers with RTSP profiles
- SLB SMTP —Traffic served by virtual servers with SMTP profiles

Within each category, you can use Filter Setting controls to filter the table based on the values of matching data:

- Date
- Time
- Proto
- Service
- Src
- Src_port
- Dst
- Dst_port
- Policy
- Action

The last column in each table includes a link to log details.

Before you begin:

- You must have Read-Write permission for Log & Report settings.

To view and filter the log:

1. Go to Log & Report > Log Access > Traffic Logs to display the traffic log.
2. Click **Filter Settings** to display the filter tools.
3. Use the tools to filter on key columns and values.
4. Click **Apply** to apply the filter and redisplay the log.

Table 142 to Table 147 list the log columns in the order in which they appear in the log.

Table 142: SLB Layer 4 and SLB TCPS logs

Column	Example	Description
date	date=2014-12-01	Log date.
time	time=07:50:36	Log time.
log_id	log_id=0102007810	Log ID.
type	type=traffic	Log type.
subtype	subtype=slb_tcps	Log subtype: slb_layer4, slb_tcps.
pri	pri=information	Log level.
vd	vd=root	Virtual domain.

Column	Example	Description
msg_id	msg_id=522030	Message ID.
duration	duration=55	Session duration.
ibytes	ibytes=138	Bytes in.
obytes	obytes=303	Bytes out.
proto	proto=6	Protocol.
service	service=tcps	Service.
src	src=31.1.1.103	Source IP address in traffic received by FortiADC.
src_port	src_port=5534	Source port.
dst	dst=21.1.1.101	Destination IP address in traffic received by FortiADC (IP address of the virtual server).
dst_port	dst_port=443	Destination port.
trans_src	trans_src=31.1.1.103	Source IP address in packet sent from FortiADC. Address might have been translated.
trans_src_port	trans_src_port=5534	Source port in packet sent from FortiADC.
trans_dst	trans_dst=21.1.1.101	Destination IP address in packet sent from FortiADC (IP address of the real server).
trans_dst_port	trans_dst_port=443	Destination port in packet sent from FortiADC.
policy	policy=L7vs	Virtual server name.
action	action=none	For most logs, action=none.
srccountry	srccountry=Reserved	Location of the source IP address.
dstcountry	dstcountry=Reserved	Location of the destination IP address.
real_server	real_server=2_2_2_10	Real server configured name.

Table 143: SLB HTTP logs

Column	Example	Description
date	date=2014-12-01	Log date.
time	time=07:50:36	Log time.
log_id	log_id=0102007810	Log ID.

Column	Example	Description
type	type=traffic	Log type.
subtype	subtype=slb_http	Log subtype: slb_http.
pri	pri=information	Log level.
vd	vd=root	Virtual domain.
msg_id	msg_id=522030	Message ID.
duration	duration=55	Session duration.
ibytes	ibytes=138	Bytes in.
obytes	obytes=303	Bytes out.
proto	proto=6	Protocol.
service	service=http	Service.
src	src=31.1.1.103	Source IP address in traffic received by FortiADC.
src_port	src_port=5534	Source port.
dst	dst=21.1.1.101	Destination IP address in traffic received by FortiADC (IP address of the virtual server).
dst_port	dst_port=443	Destination port.
trans_src	trans_src=31.1.1.103	Source IP address in packet sent from FortiADC. Address might have been translated.
trans_src_port	trans_src_port=5534	Source port in packet sent from FortiADC.
trans_dst	trans_dst=21.1.1.101	Destination IP address in packet sent from FortiADC (IP address of the real server).
trans_dst_port	trans_dst_port=443	Destination port in packet sent from FortiADC.
policy	policy=L7vs	Virtual server name.
action	action=none	For most logs, action=none.
http_method	http_method=get	HTTP method.
http_host	http_host=10.61.2.100	Host IP address.
http_agent	http_agent=curl/7.29.0	HTTP agent.
http_url=	http_url=/ip.php	Base URL.

Column	Example	Description
http_qry	http_qry=unknown	URL parameters after the base URL.
http_cookie	http_cookie=unknown	Cookie name.
http_retcode	http_retcode=200	HTTP return code.
user	user=user1	User name.
usergrp	usergrp=companyABC	User group.
auth_status	auth_status=success	Authentication success/failure.
srccountry	srccountry=Reserved	Location of the source IP address.
dstcountry	dstcountry=Reserved	Location of the destination IP address.
real_server	real_server=2_2_2_10	Real server configured name.

Table 144: SLB RADIUS log

Column	Example	Description
date	date=2014-12-01	Log date.
time	time=07:50:36	Log time.
log_id	log_id=0102007810	Log ID.
type	type=traffic	Log type.
subtype	subtype=slb_radius.	Log subtype: slb_radius.
pri	pri=information	Log level.
vd	vd=root	Virtual domain.
msg_id	msg_id=522030	Message ID.
duration	duration=55	Session duration.
ibytes	ibytes=138	Bytes in.
obytes	obytes=303	Bytes out.
proto	proto=6	Protocol.
service	service=radius	Service.
src	src=31.1.1.103	Source IP address in traffic received by FortiADC.

Column	Example	Description
src_port	src_port=5534	Source port.
dst	dst=21.1.1.101	Destination IP address in traffic received by FortiADC (IP address of the virtual server).
dst_port	dst_port=443	Destination port.
trans_src	trans_src=31.1.1.103	Source IP address in packet sent from FortiADC. Address might have been translated.
trans_src_port	trans_src_port=5534	Source port in packet sent from FortiADC.
trans_dst	trans_dst=21.1.1.101	Destination IP address in packet sent from FortiADC (IP address of the real server).
trans_dst_port	trans_dst_port=443	Destination port in packet sent from FortiADC.
policy	policy=L7vs	Virtual server name.
action	action=none	For RADIUS, action=auth or acct.
user	user=user1	RADIUS accounting username.
srccountry	srccountry=Reserved	Location of the source IP address.
dstcountry	dstcountry=Reserved	Location of the destination IP address.
real_server	real_server=2_2_2_10	Real server configured name.

Table 145: SLB RDP logs

Column	Example	Description
date	date=2016-03-18	Log date.
time	time=11:48:29	Log time.
log_id	log_id=107005800	Log ID.
type	type=traffic	Log type.
subtype	subtype=slb_rdp	Log subtype: slb_rdp.
pri	pri=information	Log level.
vd	vd=root	Virtual domain.
msg_id	msg_id=1321705	Message ID.
duration	duration=2	Session duration.

Column	Example	Description
ibytes	ibytes=92	Bytes in.
obytes	obytes=400	Bytes out.
proto	proto=6	Protocol.
service	service=http	Service.
src	src=192.168.1.1	Source IP address in traffic received by FortiADC.
src_port	src_port=37869	Source port.
dst	dst=192.168.1.142	Destination IP address in traffic received by FortiADC (IP address of the virtual server).
dst_port	dst_port=8080	Destination port.
trans_src	trans_src=2.2.2.2	Source IP address in packet sent from FortiADC. Address might have been translated.
trans_src_port	trans_src_port=58661	Source port in packet sent from FortiADC.
trans_dst	trans_dst=2.2.2.10	Destination IP address in packet sent from FortiADC (IP address of the real server).
trans_dst_port	trans_dst_port=80	Destination port in packet sent from FortiADC.
policy	policy=vs-l7	Virtual server name.
action	action=none	For most logs, action=none.
srccountry	srccountry=Reserved	Location of the source IP address.
dstcountry	dstcountry=Reserved	Location of the destination IP address.
real_server	real_server=r_22210	Real server configured name.

Table 146: SLB SIP logs

Column	Example	Description
date	date=2016-01-29	Log date.
time	time=18:06:48	Log time.
log_id	log_id=0106001134	Log ID.
type	type=traffic	Log type.
subtype	subtype=slb_sip	Log subtype: slb_sip.

Column	Example	Description
pri	pri=information	Log level.
vd	vd=root	Virtual domain.
msg_id	msg_id=154799	Message ID.
duration	duration=1	Session duration.
ibytes	ibytes=44346	Bytes in.
obytes	obytes=2.2.2.10	Bytes out.
proto	proto=6	Protocol.
service	service=http	Service.
src	src=N/A	Source IP address in traffic received by FortiADC.
src_port	src_port=43672	Source port.
dst	dst=192.168.1.142	Destination IP address in traffic received by FortiADC (IP address of the virtual server).
dst_port	dst_port=8080	Destination port.
trans_src	trans_src=2.2.2.2	Source IP address in packet sent from FortiADC. Address might have been translated.
trans_src_port	trans_src_port=80	Source port in packet sent from FortiADC.
trans_dst	trans_dst=N/A	Destination IP address in packet sent from FortiADC (IP address of the real server).
trans_dst_port	trans_dst_port=none	Destination port in packet sent from FortiADC.
policy	policy=invite	Virtual server name.
action	action=sip: bob@1.1.1.1 v2.0	Invite sent to.
sip_method	sip_method=from: alice@2.2.2.2	Invite sent from.
sip_uri	sip_uri=to: server@3.3.3.3	SIP server IP address.
sip_from	sip_from=callid:1111111	SIP call ID.
sip_to	sip_to=200	
sip_callid	sip_callid=Reserved	Reserved.
sip_retcode	sip_retcode=Reserved	Reserved.

Column	Example	Description
srccountry	srccountry=Reserved	Location of the source IP address.
dstcountry	dstcountry=Reserved	Location of the destination IP address.
real_server	real_server=2_2_2_10	Real server configured name.

Table 147: GLB log

Column	Example	Description
date	date=2014-12-01	Log date.
time	time=07:50:36	Log time.
log_id	log_id=0102007810	Log ID.
type	type=traffic	Log type.
subtype	subtype=dns	Log subtype: dns.
pri	pri=information	Log severity.
vd	vd=root	Virtual domain.
msg_id	msg_id=522030	Message ID.
proto	proto=6	Protocol.
src	src=31.1.1.103	Source IP address.
src_port	src_port=5534	Source port.
dst	dst=21.1.1.101	Destination IP address.
dst_port	dst_port=443	Destination port.
policy	policy=policy	Global load balancing policy name.
action	action=none	For most logs, action=none.
fqdn	fqdn=pool.ntp.org	FQDN from client request.
resip	resip=4.53.160.75	DNS response IP address.
srccountry	srccountry=Reserved	Location of the source IP address.
dstcountry	dstcountry=Reserved	Location of the destination IP address.

Using the script log

The Script Log table shows all the scripts used in the system.

Note: This feature is available for the SLB (server load balance) module only.

Using the aggregate log

The Aggregate Log provides an aggregated view of security logs within a selected time frame.

There are four types of aggregated security logs:

- Synflood—Traffic logged by the SYN Flood feature
- IP Reputation—Traffic logged by the IP Reputation feature
- WAF—Traffic logged by the WAF feature
- Geo—Traffic logged by the Geo IP block list feature

Each log page has two parts: left and right. The left-hand side shows the aggregated log data. Click a log entry on the left, and you'll see its details on the right.

To view an aggregated log:

1. Go to Log & Report > Log Browsing.
2. Click the **Aggregate Log** tab to display the attack log.
3. Click log type.
4. Select a time frame.
5. Click **Refresh** to apply the filter and redisplay the log.

shows the detailed information of an aggregated GEO log. The other aggregated logs show the same details.

Table 148: Details of an aggregated GEO log

Column	Example	Description
Date	2016-12-02	Log date
Time	10:27:01	Log time
Count	1	For DoS, number of timeouts sent per destination
Severity	high	Always "high" for DoS
Source	173.177.99.94	Source IP address
Destination	10.61.2.100	Destination IP address
Action	deny	Policy action

Configuring local log settings

The local log is a datastore hosted on the FortiADC system.

Typically, you use the local log to capture information about system health and system administration activities. We recommend that you use local logging during evaluation and verification of your initial deployment, and then configure remote logging to send logs to a log management repository where they can be stored long term and analyzed using preferred analytic tools.

Local log disk settings are configurable. You can select a subset of system events, traffic, and security logs.

Before you begin:

- You must have Read-Write permission for Log & Report settings.

To configure local log settings:

1. Go to Log & Report > Log Setting.
The configuration page displays the Local Log tab.
2. Complete the configuration as described in [Table 149](#).
3. Save the configuration.

Table 149: Local logging configuration

Settings	Guidelines
Status	Select to enable local logging.
File Size	Maximum disk space for a local log file. The default is 200 MB. When the current log file reaches this size, a new file is created.
Log Level	<p>Select the lowest severity to log from the following choices:</p> <ul style="list-style-type: none">• Emergency—The system has become unstable.• Alert—Immediate action is required.• Critical—Functionality is affected.• Error—An error condition exists and functionality could be affected.• Warning—Functionality might be affected.• Notification—Information about normal events.• Information—General information about system operations.• Debug—Detailed information about the system that can be used to troubleshoot unexpected behavior. <p>For example, if you select Error, the system collects logs with level Error, Critical, Alert, and Emergency. If you select Alert, the system collects logs with level Alert and Emergency.</p>

Settings	Guidelines
Disk Full	<p>Select log behavior when the maximum disk space for local logs (30% of total disk space) is reached:</p> <ul style="list-style-type: none"> • Overwrite—Continue logging. Overwrite the earliest logs. • No Log—Stop logging.
Event	Select to enable logging for events.
Event Category	<p>This option becomes available only when the Event check box is selected. In that case, select the types of events to collect in the local log:</p> <ul style="list-style-type: none"> • Configuration—Configuration changes. • Admin—Administrator actions. • System—System operations, warnings, and errors. • User—Authentication results logs. • Health Check—Health check results and client certificate validation check results. • SLB—Notifications, such as connection limit reached. • LLB—Notifications, such as bandwidth thresholds reached. • GLB—Notifications, such as the status of associated local SLB and virtual servers. • Firewall—Notifications for the "firewall" module, such as SNAT source IP pool is using all of its addresses.
Traffic	Select to enable logging for traffic processed by the load balancing modules.
Traffic Category	<p>The following options become available only when the Traffic check-box is selected. See above.</p> <ul style="list-style-type: none"> • SLB—Server Load Balancing traffic logs related to sessions and throughput. • GLB—Global Load Balancing traffic logs related to DNS requests.
Security	Select to enable logging for traffic processed by the security modules.
Security Category	<ul style="list-style-type: none"> • DoS—SYN flood protection logs. • IP Reputation—IP Reputation logs. • WAF—WAF logs. • GEO—Geo IP blocking logs.
Script	Select to enable scripting.
Script Category	SLB is selected by default and required.

Configuring syslog settings

A remote syslog server is a system provisioned specifically to collect logs for long term storage and analysis with preferred analytic tools.

Before you begin:

- You must have Read-Write permission for Log & Report settings.

To configure syslog settings:

1. Go to Log & Report > Log Setting.
2. Click the **Syslog Server** tab.
3. Click **Add** to display the configuration editor.
4. Complete the configuration as described in [Table 150](#).
5. Save the configuration.

Table 150: Syslog configuration

Settings	Guidelines
Status	Select to enable the configuration.
Address	IP address of the syslog server.
Port	Listening port number of the syslog server. Usually this is UDP port 514.
Log Level	<p>Select the lowest severity to log from the following choices:</p> <ul style="list-style-type: none">• Emergency—The system has become unstable.• Alert—Immediate action is required.• Critical—Functionality is affected.• Error—An error condition exists and functionality could be affected.• Warning—Functionality might be affected.• Notification—Information about normal events.• Information—General information about system operations.• Debug—Detailed information about the system that can be used to troubleshoot unexpected behavior. <p>For example, if you select Error, the system sends the syslog server logs with level Error, Critical, Alert, and Emergency. If you select Alert, the system collects logs with level Alert and Emergency.</p>
CSV	Send logs in CSV format. Do not use with FortiAnalyzer.
Facility	Identifier that is not used by any other device on your network when sending logs to FortiAnalyzer/syslog.

Settings	Guidelines
Event	Select to enable logging for events.
Event Category	Select the types of events to send to the syslog server: <ul style="list-style-type: none"> • Configuration—Configuration changes. • Admin—Administrator actions. • System—System operations, warnings, and errors. • User—Authentication results logs. • Health Check—Health check results and client certificate validation check results. • SLB—Notifications, such as connection limit reached. • LLB—Notifications, such as bandwidth thresholds reached. • GLB—Notifications, such as the status of associated local SLB and virtual servers. • Firewall—Notifications for the "firewall" module, such as SNAT source IP pool is using all of its addresses.
Traffic	Select to enable logging for traffic processed by the load balancing modules.
Traffic Category	<ul style="list-style-type: none"> • SLB—Server Load Balancing traffic logs related to sessions and throughput. • GLB—Global Load Balancing traffic logs related to DNS requests.
Security	Select to enable logging for traffic processed by the security modules.
Security Category	<ul style="list-style-type: none"> • DoS—SYN flood protection logs. • IP Reputation—IP Reputation logs. • WAF—WAF logs. • GEO—Geo IP blocking logs.
Script	Select to enable scripting.
Script Category	SLB is elected by default.

Configuring high speed logging

The high speed log feature is intended for deployments that require a high volume of logging activity. The logs are sent in binary format so they can be sent at a high speed. See [Appendix E: High Speed Logging Binary Format](#) for details on the structure.

The feature supports traffic logs. Event logs and security logs are not supported.

Before you begin:

- You must have Read-Write permission for Log & Report settings.

To configure high speed log settings:

1. Go to Log & Report > Log Setting.
2. Click the **High Speed Server** tab to display the configuration editor.

3. Complete the configuration as described in [Table 151](#).
4. Save the configuration.

Table 151: High speed logging configuration

Settings	Guidelines
Status	Select to enable the configuration.
Address	IP address of the syslog server.
UDP Port	Listening port number of the syslog server. Usually this is UDP port 514.
Traffic	Select to enable logging for traffic processed by the load balancing modules.
Traffic Category	The following options become available only when the Traffic check-box is selected. See above. <ul style="list-style-type: none"> • SLB—Send Server Load Balancing logs. • GLB—Send Global Load Balancing logs.
Script	Enable/disable script.
Script Category	Enable/disable server load balance log scripting.

Enabling real-time statistics

The fast statistics feature enables real-time statistics collection for fast reports. Enabled by default. You can disable fast statistics if you encounter issues.

Before you begin:

- You must have Read-Write permission for Log & Report settings.

To enable/disable real-time statistics:

1. Go to Log & Report > Log Setting.
2. Click the **Fast Stats** tab.
3. Complete the configuration as described in [Table 152](#).
4. Save the configuration.

Table 152: Fast Statistics configuration

Settings	Guidelines
Status	Enable/disable fast statistics. Enabled by default.
Traffic	Enable/disable fast statistics for traffic logs. Enabled by default.
Traffic Category	Enable/disable fast statistics for traffic categories. SLB is enabled by default.

Configuring alert email settings

You can configure alerts to be sent based on either event categories or event level (severity). See [Configuring an SMTP mail server](#) for information on how to set up the connection to the mail server.

Before you begin:

- You must have Read-Write permission for Log & Report settings.

To configure alert email settings:

1. Click Log & Report > Alert Mail.
2. Click the Setting tab.
3. Complete the configuration as described in [Table 153](#).
4. Click **Save**.

Table 153: Alert mail configuration

Settings	Guidelines
By Category	Select this option to send alerts that match the specified categories. If you do not select this option, alerts are sent based on event level (severity). See Log level below.
Category	Select the events for which alerts are sent: <ul style="list-style-type: none"> • HA • Admin • Configuration • Disk • Health Check • SSL Certificates Expire
Log Level	This option becomes available when By Category is Not enabled. Select the minimum level of severity for which alerts are sent: <ul style="list-style-type: none"> • Emergency—The system has become unstable. • Alert—Immediate action is required. • Critical—Functionality is affected. • Error—An error condition exists and functionality could be affected. • Warning—Functionality might be affected. • Notification—Information about normal events. • Information—General information about system operations. • Debug—Detailed information about the system that can be used to troubleshoot unexpected behavior. <p>For example, if you select Error, the system sends alerts with level Error, Critical, Alert, and Emergency. If you select Alert, the system sends alerts with level Alert and Emergency.</p>

Settings	Guidelines
Interval	If identical alerts are occurring continuously, select the interval at which email alerts will be sent while the event continues.
From	Sender email address used in alert email.

Configuring an alert email recipient

You can configure alerts to be sent based on either event categories or event level (severity). See [Configuring an SMTP mail server](#) for information on how to set up the connection to the mail server.

Before you begin:

- You must have Read-Write permission for Log & Report settings.

To configure an alert email recipient:

1. Click Log & Report > Alert Mail.
2. Click the **Recipient** tab.
3. Complete the configuration as described in [Table 154](#).
4. Click **Save**.

Table 154: Alert mail recipient configuration

Settings	Guidelines
Name	Recipient name to appear in alert email.
Mail To	Recipient email address.
Test Connectivity	Click this button to test the connectivity to the email address.

Configuring reports

You can configure on-demand or scheduled reports.

Before you begin:

- If you want reports to include user-defined queries, you must configure the queries before you configure the report.
- You must have Read-Write permission for Log & Report settings.

To configure a report:

1. Go to Log & Report > Report Config.
The Report tab is displayed.
2. Click **Add** to display the configuration editor.
3. Complete the configuration as described in [Table 155](#).
4. Save the configuration.

To run an on-demand report:

- In the report table, the final column for has a "run report" icon (▶). Click it.

To view a generated report:

- Go to Log & Report > Report > Overall.

Table 155: Report configuration

Settings	Guidelines
Name	Configuration name. Valid characters are A-Z, a-z, 0-9, _, and -. No spaces. You reference this name in the zone configuration (if you use forwarders). Note: After you initially save the configuration, you cannot edit the name.
On Schedule	Enable/disable reporting on schedule.
Period	Select a report period. If you select absolute or last N-hours, last N-days, or last N-weeks, additional controls are displayed for you to set these variables.
Schedule Type	Daily or on specified days.
Schedule Weekdays	If you do not schedule the report daily, specify the days on which to run it.
Schedule Hour	0-23.
Email Format	Attachment format. Only PDF is supported. If you schedule reports and set this option, the report is sent on schedule to all addresses in the Log & Report > Alert Email > Recipient list.
Email Subject	Message subject.
Email Body	Message body.
Email Attachname	Filename for attachment.
Email Compress	Enable/disable compression of the attachment.
Query List	Select queries to include in the report.

Configuring Report Queries

The predefined list of queries covers the most common administrator and stakeholder interests. It includes the following:

- SLB-Top-Policy-By-Bytes
- SLB-Top-Source-By-Bytes
- SLB-Top-Source-Country-By-Bytes
- SLB-History-Flow-By-Bytes (total traffic over time)
- LLB-Top-Link-by-Bytes
- LLB-History-Flow-By-Bytes (total traffic over time)
- DNS-Top-Policy-by-Count
- DNS-Top-Source-by-Count
- Attack-Top-Destination-For-IPReputation-By-Count
- Attack-Top-Source-For-IPReputation-By-Count
- Attack-Top-Source-Country-For-IPReputation-By-Count
- Attack-Top-Destination-For-Synflood-By-Count
- Attack-Top-Destination-For-GEO-By-Count
- Attack-Top-Source-For-GEO-By-Count
- Attack-Top-Source-Country-For-GEO-By-Count
- Attack-Top-Destination-For-WAF-By-Count
- Attack-Top-Source-For-WAF-By-Count
- Attack-Top-Source-Country-For-WAF-By-Count
- Event-Top-Admin-Login-By-Count
- Event-Top-Failed-Admin-Login-By-Count
- Event-Top-Admin-Config-By-Count

If necessary, you can create your own query configuration objects.

Before you begin:

- You must have Read-Write permission for Log & Report settings.

After you have created a query configuration object, you can select it in the report configuration.

To configure report queries:

1. Go to Log & Report > Report Config.
The Report tab is displayed.
2. Click the **Query Set** tab.
3. Click **Add** to display the configuration editor.
4. Complete the configuration as described in [Table 156](#).
5. Save the configuration.

Table 156: Query configuration

Settings	Guidelines
Name	<p>Configuration name. Valid characters are A-Z, a-z, 0-9, _, and -. No spaces. You reference this name in the zone configuration (if you use forwarders).</p> <p>Note: After you initially save the configuration, you cannot edit the name.</p>

Settings	Guidelines
Module	<ul style="list-style-type: none"> • SLB • LLB • DNS • Attack • Event
SLB	
Traffic Sort Type	<ul style="list-style-type: none"> • sessions • bytes
SLB Subtype	<ul style="list-style-type: none"> • top_policy (virtual server) • top_source • top_source_country • slb_history_flow (total traffic over time)
LLB	
Traffic Sort Type	<ul style="list-style-type: none"> • sessions • bytes
LLB Subtype	<ul style="list-style-type: none"> • top_link • slb_history_flow (total traffic over time)
DNS	
DNS Sort Type	Only count is applicable.
DNS Subtype	<ul style="list-style-type: none"> • Top_Policy • top_source
Attack	
Attack Sort Type	Only count is applicable.

Settings	Guidelines
Attack Subtype	<ul style="list-style-type: none"> • top_destip_for_geo • top_destip_for_ipreputation • top_destip_for_sysflood • top_destip_for_waf • top_source_country_for_geo • top_source_country_for_ipreputation • top_source_country_for_waf • top_source_for_geo • top_source_for_ipreputation • top_source_for_waf
Event	
Event Sort Type	Only count is applicable.
Event Subtype	<ul style="list-style-type: none"> • top_admin_login • top_failed_admin_login • top_admin_config

Configuring fast reports

Fast reports are real time statistics displayed on the Dashboard > Data Analytics page.

Before you begin:

- You must have Read-Write permission for Log & Report settings.

After you have created a query configuration object, you can select it in the report configuration.

To configure a fast report:

1. Go to Dashboard > Data Analytics.
2. Click **Add Widget** to display the configuration editor.
3. Complete the configuration as described in [Table 157](#).
4. Save the configuration.

Table 157: Fast report configuration

Settings	Guidelines
Name	<p>Configuration name. Valid characters are A-Z, a-z, 0-9, _, and -. No spaces. You reference this name in the zone configuration (if you use forwarders).</p> <p>Note: After you initially save the configuration, you cannot edit the name.</p>

Settings	Guidelines
Module	SLB. Note: For the current release, SLB is the only option and is selected by default.
SLB SubType	Select an option from the list menu: <ul style="list-style-type: none">• Top Src• Top Dest• Top Browser• Top OS• Top Dev• Top Domain• Top URL• Top Referrer• Top Source Country• Top Session
History Chart	Enable/Disable.
Time Range	Select an option from the list menu: <ul style="list-style-type: none">• 10MINS• 1HOUR• 1DAY• 1WEEK• 1MONTH
Data Type	Select either of the following: <ul style="list-style-type: none">• Bandwidth• Session

Viewing reports

FortiADC provides the following reports:

- [Overall](#)
- [Server Load Balance](#)
- [Link Load Balance](#)
- [Global Load Balance](#)
- [Security](#)

Viewing the Overall report

The Overall tab has a table that lists generated reports. The last column in the table has icon links you can use to delete the generated report, preview it, or download it.

Before you begin:

- You must have generated reports from the Log & Report > Report Config > Report page.
- You must have Read-Write permission for Log & Report settings.

To navigate to the report table, go to Log & Report > Report > Overall.

Figure 58: Overall tab: Table of reports



















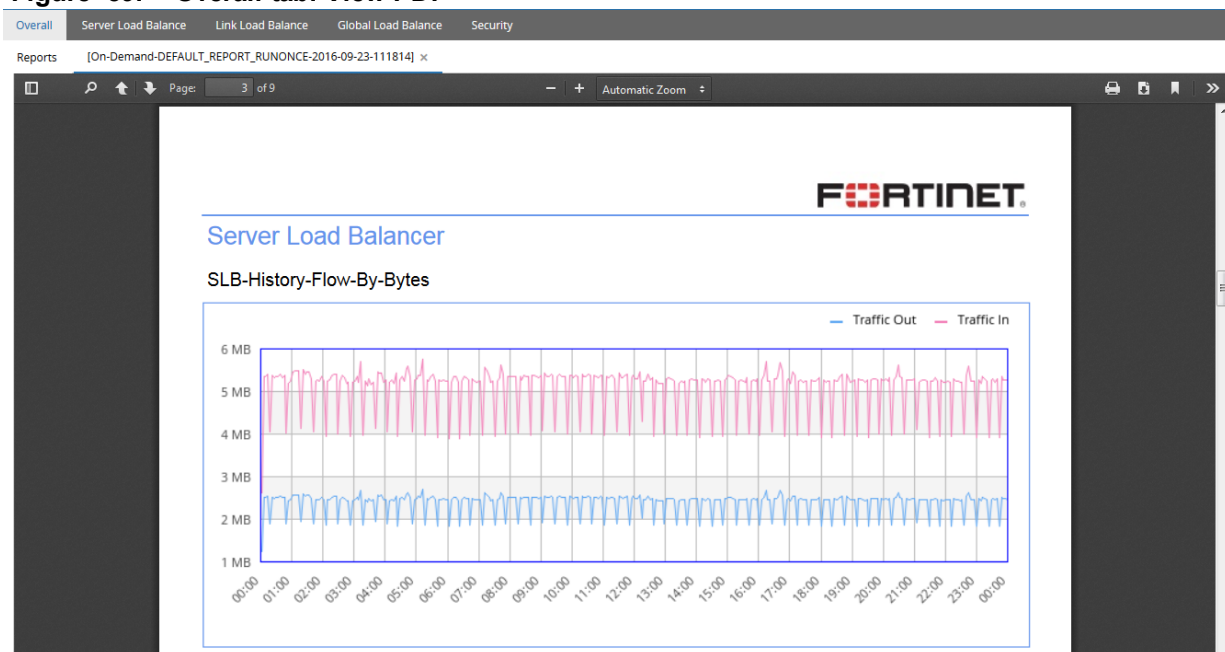
Overall Server Load Balance Link Load Balance Global Load Balance Security					
Reports Refresh					
Report File	Status	Period	Create Time	Size	
On-Demand-DEFAULT_REPORT_RUNONCE-2016-05-23-073719	Done	2016-05-22-00:00:00-2016-05-22-23:59:59	2016-05-23 07:38:29	101.76 KB	  
On-Demand-DEFAULT_REPORT_RUNONCE-2016-05-02-232559	Done	2016-05-01-00:00:00-2016-05-01-23:59:59	2016-05-02 14:26:41	99.28 KB	  
On-Demand-DEFAULT_REPORT_RUNONCE-2016-04-13-093113	Done	2016-04-12-00:00:00-2016-04-12-23:59:59	2016-04-13 00:31:54	116.69 KB	  
On-Demand-DEFAULT_REPORT_RUNONCE-2016-02-18-205325	Done	2016-02-17-00:00:00-2016-02-17-23:59:59	2016-02-18 11:54:07	109.39 KB	  
On-Demand-DEFAULT_REPORT_RUNONCE-2016-02-18-205327	Done	2016-02-17-00:00:00-2016-02-17-23:59:59	2016-02-18 11:54:45	109.39 KB	  
On-Demand-DEFAULT_REPORT_RUNONCE-2015-11-09-224341	Done	2015-11-04-00:00:00-2015-11-10-00:00:00	0	0	  

Figure 59: Overall tab: View PDF



Viewing the Server Load Balance report

The Server Load Balance report is a graph of throughput for the SLB virtual servers with the top throughput.

Before you begin:

- You must have Read-Write permission for Log & Report settings.

To generate a report:

- Go to Log & Report > Report.
- Click the **Server Load Balance** tab.
- Specify a time range for the report.
- Select the number of virtual servers to include in the graph.
- Select whether to use inbound traffic, outbound traffic, or number of sessions as the measure of throughput.
- Click **Search** to generate the query.

A summary table is displayed. See [Figure 60](#).

- Drill into the details for a virtual server by clicking the query links:

Table 158: Icons on Server Load Balance report page

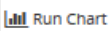



	Generate a report that shows the trend in throughput over time. See Figure 61 .
	Generate a report with the top client data. You can query by Source IP address, OS Type, or Browser Type. See Figure 62 .
	Note: OS Type and Browser Type queries apply to HTTP servers only.
	Generate a report with the top destination URLs. See Figure 63 .
	Note: Top URLs applies to HTTP servers only.
	Drill into queries for the real servers or content routes that belong to this virtual server. See Figure 64 .

Figure 60: Server Load Balance summary list

Overall

Server Load Balance

Link Load Balance

Global Load Balance

Security

Start Date Time: 2016-10-06-00:00 - End Date Time: 2016-10-13-00:00

Top 5

By inbytes

Search

Virtual Server By inbytes Top 5

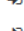


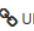

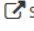


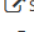
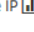

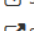
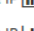




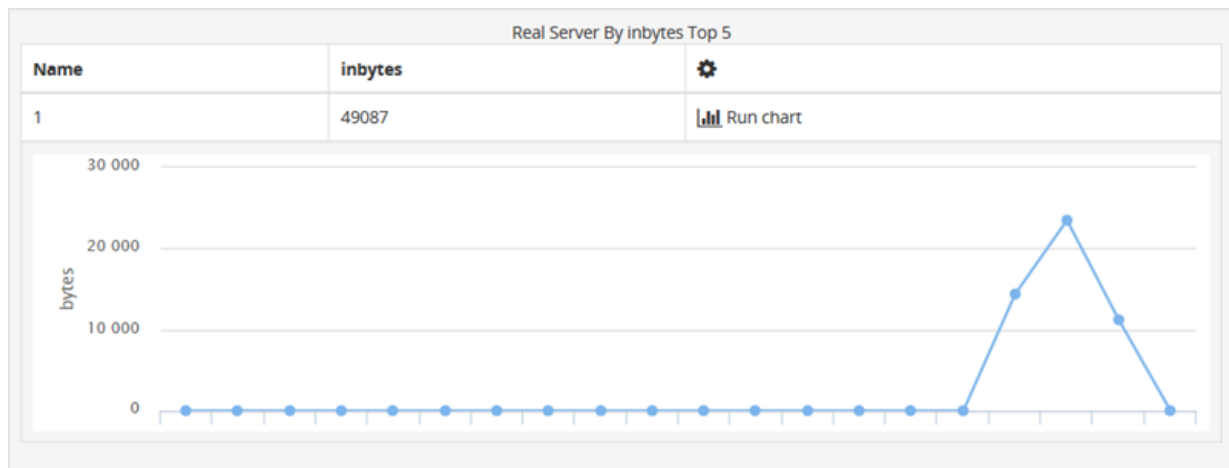
Name	Inbytes	Content Route	Pool Name	Real Server	
vs1	203079410		server1		 Client  URL  Run Chart
vs2	43985576		server1		 Source IP  Run Chart
rad-VS	0		server1		 Source IP  Run Chart
ssh-loadbalancing	0		ssh-Server3		 Source IP  Run Chart
test-ssl	0		server1		 Source IP  Run Chart

Figure 61: Virtual Server throughput (Run Chart)**Figure 62: Virtual Server client reports**

Source IP By inbytes Top 5		Source IP	OS Type	Browser Type
IP	Inbytes			
192.168.1.1	47611			
192.168.1.100	1476			
OS Type By count Top 5				
Name	count			
Linux	380			
Unknown	1			
Browser Type By count Top 5				
Name	count			
Unknown	371			
Firefox/39.0	8			
Chrome/44.0.2403.89	2			

Figure 63: Virtual server URL report

waf	49087		2.2.2.10	➔	Client URL Run Chart
URL By count Top 5					
Name	count				
/bigdata	369				
/	8				
/favicon.ico	3				
unknown	1				

Figure 64: Real Server throughput (Run Chart)

Viewing the Link Load Balance report

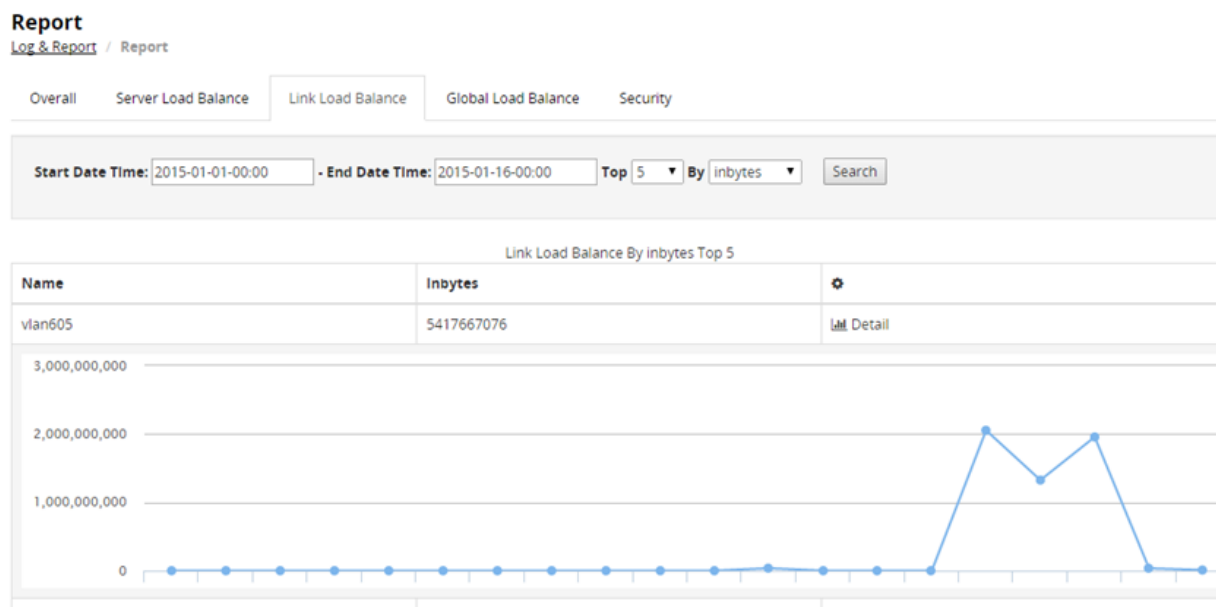
The Link Load Balance report is a graph of throughput for the LLB gateways with the top throughput.

Before you begin:

- You must have Read-Write permission for Log & Report settings.

To generate a report:

- Go to Log & Report > Report.
- Click the **Link Load Balance** tab.
- Specify a time range for the report.
- Select whether to use inbound traffic, outbound traffic, or number of sessions as the measure of throughput.
- Select the number of gateways to include in the graph.
- Click **Search** to generate the query.
A summary table is displayed.
- Click **Details** to display the report graph.

Figure 65: Link Load Balance report

Viewing the Global Load Balance report

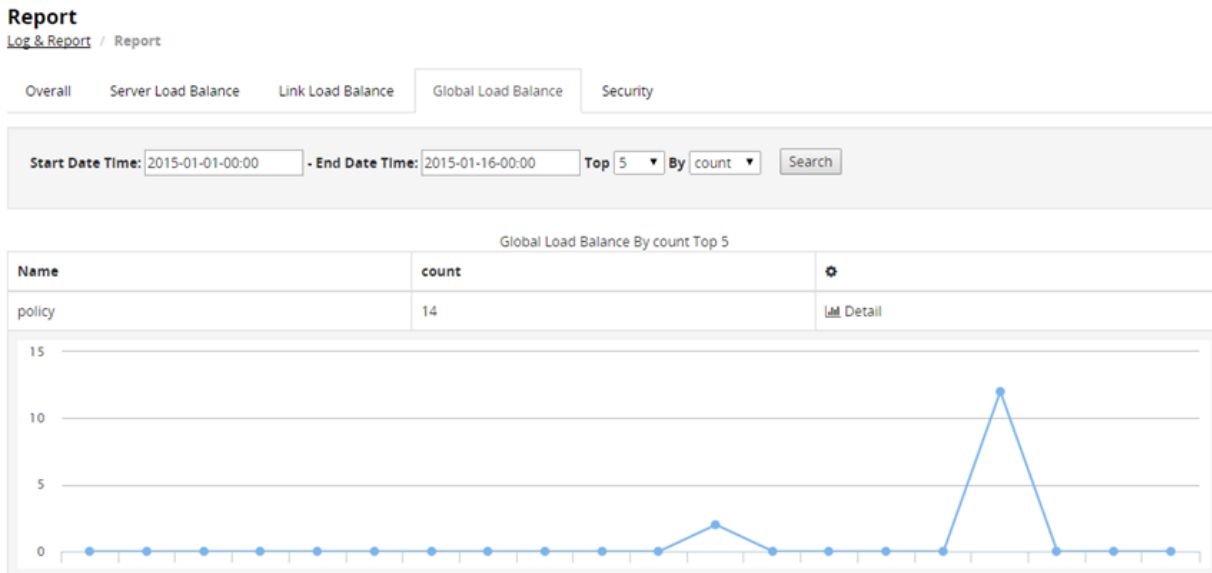
The Global Load Balance report is a graph of traffic matching global load balancing policies.

Before you begin:

- You must have Read-Write permission for Log & Report settings.

To generate a report:

1. Go to Log & Report > Report.
2. Click the **Global Load Balance** tab.
3. Specify a time range for the report.
4. Select the number of policies to include in the graph.
The report is added to the report table.
5. Click **Search** to generate the query.
A summary table is displayed.
6. Click **Details** to display the report graph.

Figure 66: Global load balance report

Viewing the Security report

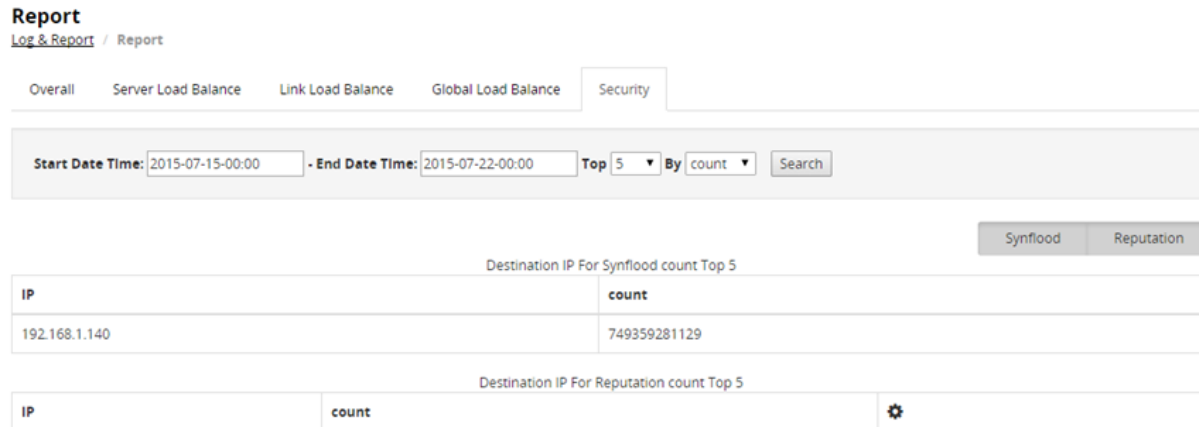
The Security report is a count of traffic matching security policies.

Before you begin:

- You must have Read-Write permission for Log & Report settings.

To generate a report:

- Go to Log & Report > Report.
- Click the **Security** tab.
- Specify a time range for the report.
- Select the number of destination IP addresses to include in the graph.
- Click **Search** to generate the query.
A summary table is displayed.
- Click **Details** to display the report graph.

Figure 67: Security report

Display logs via CLI

FortiADC allows you to display logs using the CLI, with filtering functions.

```
FortiADC-VM # execute log
delete-file      delete-file
delete-type      delete-type
display          display the log message
filter           set filter for log browsing
list-type        list-type
rebuild-db       rebuild-db
```

```
FortiADC-VM # execute log display

<startline integer >=0 >    show log from startline

FortiADC-VM # execute log filter

<type|subtype|field|clear|show>    set ,clear,show filters
```

Where:

- type <event|traffic|attack>
- subtype <subtype_value> ex:slb_http
- field <field_name> <field_value_list>

Chapter 14: High Availability Deployments

This chapter includes the following topics:

- [HA feature overview](#)
- [HA system requirements](#)
- [HA configuration synchronization](#)
- [Configuring HA settings](#)
- [Monitoring an HA cluster](#)
- [Updating firmware for an HA cluster](#)
- [Deploying an active-passive cluster](#)
- [Deploying an active-active cluster](#)
- [Deploying an active-active-VRRP cluster](#)

HA feature overview

FortiADC appliances can be deployed as standalone units or as high availability (HA) clusters.

A *cluster* is two or more nodes. A *node* is an instance of the appliance/system. In a cluster, one node is the *primary node*, also called the *master node*. The other members of the cluster are *secondary nodes*, also called *slave nodes*.

The primary node has a special role. It has a one-to-many relationship with member nodes. Both configuration updates and software updates are initiated by the primary node and pushed to member nodes.

The system selects the primary node based on the following criteria:

- Link health (if monitor ports links are down, the node is considered down)
- Remote IP monitor health check results
- Override setting (prefers priority to uptime)
- Most available ports
- Highest uptime value
- Lowest device priority number (1 has greater priority than 2)
- Highest-sorting serial number—Serial numbers are sorted by comparing each character from left to right, where 9 and z are the greatest values. The system gives preference to higher values over lower values.

HA solutions depend on two types of communication among cluster members:

- Synchronization—During initialization, the primary node pushes its configuration (with noted exceptions) to member nodes. After initialization has completed, the nodes synchronize their session tables.
- Heartbeats—A cluster node indicates to other nodes in the cluster that it is up and available. The absence of heartbeat traffic indicates the node is not up and is unavailable.

There are three types of HA clusters:

- Active-Passive—Only the primary node is active, so it is the only node that receives traffic from adjacent routers. Typically, there is one other node that is in standby mode. It assumes active status if the primary node undergoes

maintenance or otherwise becomes unavailable.

- **Active-Active**—All nodes receive traffic. Active-Active deployments support load balancing and failover among up to eight cluster members.
- **Active-Active-VRRP** —FortiADC's Active-Active-VRRP mode uses a VRRP-like protocol, and can function in both HA Active-Passive mode and HA Active-Active mode, depending on the number of traffic groups used in the configuration. When only one traffic group is used, it actually functions in Active-Passive mode; when two or more traffic groups are used, it works in Active-Active mode.

In an Active-Passive cluster, only the management IP address for the primary node is active. In an active-passive cluster, you can log into a node only when it has primary node status and its IP address is active. To access the user interface of an appliance in standby status (the active-passive slave), you must use a console port connection.

In an Active-Active cluster, the IP addresses for all interfaces are unique, including the management interface. When the appliance is in standalone mode, the physical port IP address is active; when it is in HA mode, the address assigned to it in the HA node IP list address is active. You can log into any node using the active IP address for its management port.

In an Active-Active-VRRP cluster, FortiADC uses hbdev for members status communication. It also allows you to configure sync+session, persistence sync, and image sync functions via hbdev and dataport, which is essentially the same as the HA-AA/AP mode. Note that FortiADC is unable to communicate with third-party VRRP devices because it actually doesn't use the VRRP protocol at all.

Tip: You can use the `execute ha manage` command to log into the console of a member node. See the CLI reference.

Figure 68 shows an *active-passive* cluster in a single network path. In an active-passive cluster, the primary node is the active node that handles all traffic. In the event that the primary node experiences hardware failure or system maintenance, *failover* takes place. In failover, the standby node becomes the primary node and processes the traffic that is forwarded along the network path. The new primary node sends gratuitous ARP to notify the network to direct traffic for the virtual MAC addresses (vMAC) to its network interfaces. It takes the IP addresses of the unresponsive node.

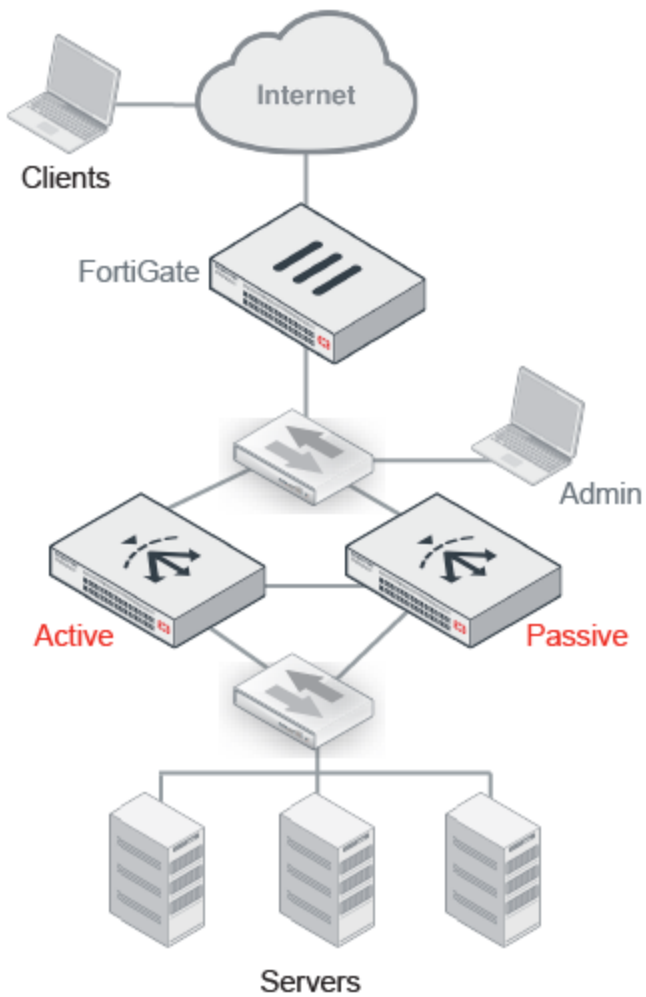
Figure 68: Basic active-passive cluster

Figure 69 shows an active-passive cluster in a *redundant path*. A topology like this is a best practice because it is fully redundant, with no single point of failure. If the gateway, load balancer, or switch were to fail, the failover path is chosen.



Figure 69: Redundant path active-passive cluster

Figure 70 shows an *active-active* cluster. An active-active cluster supports load-balancing and failover among up to eight member nodes. The routers on either side of the cluster must be configured to use equal cost multipath (ECMP) to distribute traffic to the FortiADC cluster nodes. All nodes actively receive and forward traffic.

The primary node has a special role. It handles all FTP and firewall traffic, and it acts as the failover node for all of the other nodes in the cluster.

The failover mechanism is the same as an active-passive deployment, with the primary node acting as the standby node for all other cluster members. If a member node fails, the primary node takes the IP addresses of the unresponsive node and notifies the network via ARP to redirect traffic for that vMAC to its own network interfaces. For example, in Figure 70, node1 is the primary node. If node2 were to fail, its traffic would failover to node1. If node3 were to fail, its traffic would also failover to node1. If the primary node were to fail, a new primary node would be elected, and it would function as the master in all respects, including its role as the new standby node for failover from all other cluster members.

Figure 70: Basic active-active cluster

HA system requirements

- Appliances must have the same hardware model and same firmware version.
- Redundant network topology: if an active node fails, physical network cabling and routes must be able to redirect traffic to the other member nodes.
- At least one physical port on both HA appliances to be used for heartbeat and data traffic between cluster members. For active-passive failover pairs, you can connect the ports directly via a crossover cable. For active-active clusters with more than two members, you can connect the nodes via the same Layer 2 switch.
- Heartbeat and synchronization traffic between cluster nodes occur over the physical network ports that you designate. If switches are used to connect the nodes, the interfaces must be reachable by Layer 2 multicast.
- Each appliance must be licensed. If using FortiADC-VM, the license must be paid; trial licenses will not function.



FortiADC-VM supports HA. However, if you do not want to use the native HA, you can use your hypervisor or VM environment manager to install your virtual appliances over a hardware cluster to improve availability. For example, VMware clusters can use vMotion or VMware HA.

HA configuration synchronization

Normally in an HA configuration, the master node pushes most of its configuration to the other member nodes. This is known as HA configuration synchronization. If automatic synchronization is enabled, synchronization occurs automatically when an appliance joins the cluster, and it repeats every 30 seconds thereafter. If synchronization is not enabled, you must initiate synchronization manually.

HA configuration synchronization includes:

- Core CLI-style configuration file (fadc_system.conf)
- X.509 certificates, certificate signing request files (CSR), and private keys
- Layer-7 virtual server error message files
- Layer-4 TCP connection state, Layer-4 persistence table, and Layer-7 persistence table (Source Address Persistence table only)
- Health check status (active-passive deployments only)

For most settings, you configure only the primary node, and its settings are pushed to other members.

[Table 159](#) summarizes the configuration settings that are not synchronized. All other settings are synchronized.

Table 159: HA settings that are not synchronized

Setting	Explanation
Hostname	The hostnames are not synchronized to enable you to use unique names.
SNMP system information	Each member node has its own SNMP system information so that you can maintain accurate, separate data in SNMP collections. However, the network interfaces of a standby node are not active, so they cannot be actively monitored with SNMP.
RAID level	RAID settings are hardware-dependent and determined at boot time by looking at the drives (for software RAID) or the controller (hardware RAID), and are not stored in the system configuration. Therefore, they are not synchronized.

Setting	Explanation
HA settings	<p>Most of the HA configuration is not synchronized in order to support HA system operations. In particular:</p> <ul style="list-style-type: none"> • Priority and Override settings—These settings are used to elect a primary node, so they are not synchronized to enable differentiation. • Group ID—Nodes with the same Group ID join a cluster. The setting precedes and determines group membership, so it is set manually. • HA mode—Many administrators prefer to be able to switch the primary node from an HA mode to standalone mode without the other nodes following suit, or to switch a secondary node to standalone mode and have that setting not overwritten by periodic synchronization, so the HA mode setting is not pushed from the primary node to the member nodes. • Node list and Local Node ID—These settings are for active-active mode only. They identify a node uniquely within an active-active load balancing group, so they are not synchronized to enable differentiation.

In addition to HA settings, the following data is *not* synchronized either:

- Log messages—These describe events that happened on a specific appliance. After a fail-over, you might notice that there is a gap in the original active appliance's log files that corresponds to the period of its down time. Log messages created during the time when the standby was acting as the active appliance (if you have configured local log storage) are stored there, on the original standby appliance.
- Generated reports—Like the log messages that they are based upon, reports also describe events that happened on that specific appliance. As such, report settings are synchronized, but report output is not.

You can view the status of cluster members from the dashboard of the primary node. You might need to log into the system for a non-primary member node in the following situations:

- To configure settings that are not synchronized.
- To view log messages recorded about the member node itself on its own hard disk.
- To view traffic reports for traffic processed by the member node.

Configuring HA settings

Note: Currently, FortiADC only supports HA configurations for IPv4 address mode; HA is not supported on IPv6.

Before you begin:

- You must have Read-Write permission to items in the System category.

To configure HA settings:

1. Go to System > High Availability.
2. Complete the configuration as described in [Table 160](#).
3. Save the configuration.

After you have saved the configuration, cluster members begin to send heartbeat traffic to each other. Members with the same Group ID join the cluster. They send synchronization traffic through their data links.

Table 160: High availability configuration

Settings	Guidelines
Operation Mode	<ul style="list-style-type: none"> • Standalone • Cluster
Cluster Mode	<ul style="list-style-type: none"> • Active-Passive • Active-Active • Active-Active-VRRP
Group Name	Name to identify the HA cluster if you have more than one. This setting is optional, and does not affect HA function. The maximum length is 63 characters.
Group ID	Number that identifies the HA cluster. Nodes with the same group ID join the cluster. If you have more than one HA cluster on the same network, each cluster must have a different group ID. The group ID is used in the virtual MAC address that is sent in broadcast ARP messages. The valid range is 0 to 31. The default value is 0.
Priority	<p>Number indicating priority of the member node when electing the cluster primary node. This setting is optional. The smaller the number, the higher the priority. The default is 5. The valid range is 0 to 9.</p> <p>Note: By default, uptime is more important than this setting unless Override is enabled.</p>
Config Priority	<p>The default value is 100, but you can specify any numeric value ranging from 0 to 255.</p> <p>Note: FortiADC 4.7.x has introduced a new parameter called config-priority for HA configuration. It allows you to determine which configuration the system uses when synchronizing the configuration between the HA nodes. Therefore, upon upgrading to FortiADC 4.7.x, it is highly recommended that you use this option to manually set different HA configuration priority values on the nodes. Otherwise, you'll have no control over the system's master-slave configuration sync behavior. When the configuration priority values are identical on both nodes (whether by default or by configuration), the system uses the configuration of the appliance with the larger serial number to override that of the appliance with the smaller serial number. When the configuration priority values on the nodes are different, the configuration of the appliance with the lower configuration priority will prevail.</p>
Override	Enable to make Device Priority a more important factor than uptime when selecting the primary node.

Settings	Guidelines
Heartbeat Interval	<p>Number of 100-millisecond intervals at which heartbeat packets are sent. This is also the interval at which a node expects to receive heartbeat packets. This part of the configuration is pushed from the primary node to member nodes. The default is 2. The valid range is 1 to 20 (that is, between 100 and 2,000 milliseconds).</p> <p>Note: Although this setting is pushed from the primary node to member nodes, you should initially configure all nodes with the same Detection Interval to prevent inadvertent failover from occurring before the initial synchronization.</p>
Lost Heartbeat Threshold	<p>Number of times a node retries the heartbeat and waits to receive HA heartbeat packets from the other nodes before concluding the other node is down. This part of the configuration is pushed from the primary node to member nodes. Normally, you do not need to change this setting. Exceptions include:</p> <ul style="list-style-type: none">• Increase the failure detection threshold if a failure is detected when none has actually occurred. For example, in an active-passive deployment, if the primary node is very busy during peak traffic times, it might not respond to heartbeat packets in time, and a standby node might assume that the primary node has failed.• Decrease the failure detection threshold or detection interval if administrators and HTTP clients have to wait too long before being able to connect through the primary node, resulting in noticeable down time. <p>The valid range is from 1 to 60.</p> <p>Note: Although this setting is pushed from the primary node to member nodes, you should initially configure all nodes with the same HB Lost Threshold to prevent inadvertent failover from occurring before the initial synchronization.</p>

Settings	Guidelines
ARP Times	<p>Number of times that the cluster member broadcasts extra address resolution protocol (ARP) packets when it takes on the primary role. (Even though a new NIC has not actually been connected to the network, the member does this to notify the network that a new physical port has become associated with the IP address and virtual MAC of the HA cluster.) This is sometimes called “using gratuitous ARP packets to train the network,” and can occur when the primary node is starting up, or during a failover. Also configure ARP Packet Interval.</p> <p>Normally, you do not need to change this setting. Exceptions include:</p> <ul style="list-style-type: none"> • Increase the number of times the primary node sends gratuitous ARP packets if an active-passive cluster takes a long time to fail over or to train the network. Sending more gratuitous ARP packets may help the failover to happen faster. • Decrease the number of times the primary node sends gratuitous ARP packets if the cluster has a large number of VLAN interfaces and virtual domains. Because gratuitous ARP packets are broadcast, sending them might generate a large amount of network traffic. As long as the active-passive cluster fails over successfully, you can reduce the number of times gratuitous ARP packets are sent to reduce the amount of traffic produced by a failover. <p>The valid range is 1 to 60. The default is 5.</p>
ARP Interval	<p>Number of seconds to wait between each broadcast of ARP packets. Normally, you do not need to change this setting. Exceptions include:</p> <ul style="list-style-type: none"> • Decrease the interval if an active-passive cluster takes a long time to fail over or to train the network. Sending ARP packets more frequently may help the failover to happen faster. • Increase the interval if the cluster has a large number of VLAN interfaces and virtual domains. Because gratuitous ARP packets are broadcast, sending them might generate a large amount of network traffic. As long as the active-passive cluster fails over successfully, you can increase the interval between when gratuitous ARP packets are sent to reduce the rate of traffic produced by a failover. <p>The valid range is from 1 to 20. The default is 6 seconds.</p>

Settings	Guidelines
Layer 7 Persistence Synchronization	<p>Enable to synchronize Layer 7 session data used for persistence to backend servers.</p> <p>When enabled, the Source Address Persistence table is synchronized between HA members.</p> <p>When not enabled, a node that receives traffic due to failover would not know that a session had been created already, so it will be treated as a new session.</p> <p>Synchronization of the persistence table is not required for cookie-based or hash-based persistence methods to get the desired result. Client traffic will be routed to the same backend server.</p> <p>Synchronization of the persistence table is not possible for SSL session ID. When the session via the first node is terminated, the client must re-establish an SSL connection via the second node. When a client requests a new SSL connection with an SSL server, the initial TCP connection has an SSL Session ID of 0. This zero value tells the server that it needs to set up a new SSL session and to generate an SSL Session ID. The server sends the new SSL Session ID in its response to the client as part of the SSL handshake.</p>
Layer 4 Persistence Synchronization	<p>Enable to synchronize Layer 4 session data used for persistence to backend servers.</p> <p>When enabled, the Source Address Persistence table is synchronized between HA members. When not enabled, a node that receives traffic because of load balancing or failover would not know that a session had been created already, so it will be treated as a new session.</p> <p>Synchronization of the persistence table is not required for hash-based persistence methods to get the desired result. Client traffic will be routed to the same backend server.</p>
Layer 4 Connection Synchronization	<p>Enable to synchronize Layer 4 connection state data.</p> <p>When enabled, the TCP session table is synchronized. If subsequent traffic for the connection is distributed through a different cluster node because of failover, the TCP sessions can resume without interruption.</p> <p>When not enabled, a node that receives traffic because of failover would not know that a session had been created already, and the client will be required to re-initialize the connection.</p>
Auto Config Sync	<p>Enable/disable automatic configuration synchronization. When enabled, synchronization occurs immediately when an appliance joins the cluster, and thereafter every 30 seconds. Disable if you prefer to manage synchronization manually.</p>
Active-Active Settings	
Node List	<p>Select the node IDs for the nodes in the cluster. An active-active cluster can have up to eight members.</p>

Settings	Guidelines
Local Node	A number that uniquely identifies the member within the cluster. The valid range is 0-7. In an active-active deployment, this number is used in the virtual MAC address that is sent in ARP responses. In an active-passive deployment, this number is not used.
Port Monitor	
Monitor	<p>One or more network interfaces that correlate with a physical link. These ports will be monitored for link failure. Port monitoring (also called interface monitoring) monitors physical network ports to verify that they are functioning properly and linked to their networks. You can monitor physical interfaces and 802.3ad aggregated interfaces.</p> <p>Note: To prevent an unintentional failover, do not configure port monitoring <i>until</i> you configure HA on all appliances and have plugged in the cables to link the physical network ports that will be monitored.</p>
Heartbeat	<p>Set the network interface to be used for heartbeat packets. You can configure one or two heartbeat ports.</p> <p>Use the same port number for all cluster members. For example, if you select port3 on the primary node, select port3 as the heartbeat interface on the other member nodes.</p> <p>Note: If a switch is used to connect the heartbeat interfaces, the heartbeat interfaces must be reachable by Layer 2 multicast</p> <p>Note: In order for the HA feature to work properly on VMware VMs, you MUST ensure that the vSphere vSwitch VLAN used for the heartbeat interface is able to accept MAC Address Changes and Forced Transmits..For more information, refer to FortiADC-VM™ Install Guide — D-Series.</p>
Data	Set the network interface to be used for data synchronization among cluster nodes. You can configure up to two data ports. If one data port fails, its traffic fails over to the next data port. If all data ports fail, data synchronization traffic fails over to the heartbeat port. If you do not configure a data port, the heartbeat port is used for synchronization. Use the same port numbers for all cluster members. For example, if you select port3 on the primary node, select port3 as the data port interface on the other member nodes.
Remote IP Monitor	
Monitor Enable	Enable/disable active monitoring of remote beacon IP addresses to determine if the network path is available.
Failover Threshold	Number of consecutive times that the remote IP address is unreachable that indicates failure. The default is 5. The valid range is 1-300.
Failover Hold Time	If failover occurs due to a remote IP monitor test, and this node's role changes (to master or slave), it cannot change again until the holdtime elapses. Holdtime can be used to prevent looping. The default holdtime is 120 seconds. The valid range is 60-86400.

Settings	Guidelines
Remote IP Monitor List	
Name	Configuration name. Valid characters are A-Z, a-z, 0-9, _, and -. No spaces. You reference this name in the virtual server configuration. Note: After you initially save the configuration, you cannot edit the name.
Remote Address	Remote address to ping.
Source Port	Interface to send the health check ping.
Health Check Interval	Seconds between each health check. Should be more than the timeout to prevent overlapping health checks. The default is 10.
Health Check Timeout	Seconds to wait for a reply before assuming that the health check has failed. The default is 5.
Health Check Retry	Number of retries to confirm up or down. The default is 3 retries. The valid range is 1-10.

Monitoring an HA cluster

You can view HA status from the system dashboard. Go to System > Dashboard and click the HA Status tab.

Figure 71: HA Status page

StatusData AnalyticsServer Load BalanceLink Load BalanceGlobal Load BalanceHA StatusSession Monitoring

HA StatusTraffic Status

Mode:standalone

State:Standalone

Config Sync:N/A

Serial Number:FADV0000000TRIAL

Node ID:0

IP Address:169.254.107.191

Last Changed Time:Wed Dec 31 16:00:00 1969

Last Changed Reason:

Sync Statistics

Sync Pkts	Sent	Received
L4 Session and Persistence Sync Pkts	0	0
L7 Persistence Sync Pkts	0	0

Device Management Errors

Duplicate Node ID:0


Version Mismatch:0

You can also use log messages, alert emails, and SNMP to monitor HA events, such as when failover has occurred. The system logs HA node status changes as follows:

- When HA is initialized: `HA device Init`
- When a member joins a group: `Member (FAD2HD3A12000003) join to the HA group`
- When the HA configuration is changed from standalone to an active-passive or active-active cluster mode: `HA device into Slave mode`

The following figure shows FortiADC HA event objects in an SNMP manager.

Figure 72: FortiADC HA event objects in an SNMP manager



Updating firmware for an HA cluster

You can upgrade firmware on all nodes in a cluster from the primary node.

The following process occurs when you perform the HA upgrade procedure:

1. The primary node pushes the firmware image to the member nodes.
2. The primary node notifies the member nodes of the upgrade, and it takes their user traffic during the upgrade.
3. The upgrade command is run on the member nodes, the systems are rebooted, and the member nodes send the primary node an acknowledgment that upgrade has been completed.
4. The upgrade command is run on the primary node, and it reboots. When the system is rebooting, a member node assumes primary status, and the traffic fails over from the former primary node to the new primary node.

After the upgrade process is completed, the system determines whether the original node becomes the primary node, according to the HA Override setting:

- If Override is *enabled*, the cluster considers the Device Priority setting. Both nodes usually make a second failover in order to resume their original roles.
- If Override is *disabled*, the cluster considers uptime first. The original primary node will have a smaller uptime due to the order of reboots during the firmware upgrade. Therefore it will *not* resume its active role; instead, the node with the greatest uptime will remain the new primary node. A second failover will *not* occur.

Reboot times vary by the appliance model, and also by differences between the original firmware version and the firmware version you are installing.

The administrator procedure for an HA cluster is similar to the procedure for installing firmware on a standalone appliance. To ensure minimal interruption of service to clients, use the following steps. The same procedure applies to both active-active and active-passive clusters.



If *downgrading* to a previous version, do *not* use this procedure. The HA daemon on a member node might detect that the primary node has older firmware, and attempt to upgrade it to bring it into sync, undoing your downgrade.

Instead, switch out of HA, downgrade each node individually, then switch them back into HA mode.

Before you begin:

- Download the firmware file from the Fortinet Customer Service & Support website: <https://support.fortinet.com/>
- Read the release notes for the version you plan to install.
- Back up your configuration before beginning this procedure. Reverting to an earlier firmware version could reset settings that are not compatible with the new firmware.
- You must have super user permission (user **admin**) to upgrade firmware.
- Verify that the cluster node members are powered on and available on *all* of the network interfaces that you have configured. If required ports are not available, HA port monitoring could inadvertently trigger an additional failover, resulting in traffic interruption during the firmware update.

To update the firmware for an HA cluster:

1. Log into the web UI of the *primary* node as the `admin` administrator.
2. Go to System > Settings.
3. Click the **Maintenance** tab.
4. Scroll to the Upgrade section.
5. Click **Choose File** to locate and select the file.
6. Enable the HA Sync.
7. Click ⓘ to upload the firmware and start the upgrade process.

After the new firmware has been installed, the system reboots.

When you update software, you are also updating the web UI. To ensure the web UI displays the updated pages correctly:



- Clear your browser cache.
- Refresh the page.

In most environments, press Ctrl-F5 to force the browser to get a new copy of the content from the web application. See the Wikipedia article on browser caching issues for a summary of tips for many environments:

https://en.wikipedia.org/wiki/Wikipedia:Bypass_your_cache.

Deploying an active-passive cluster

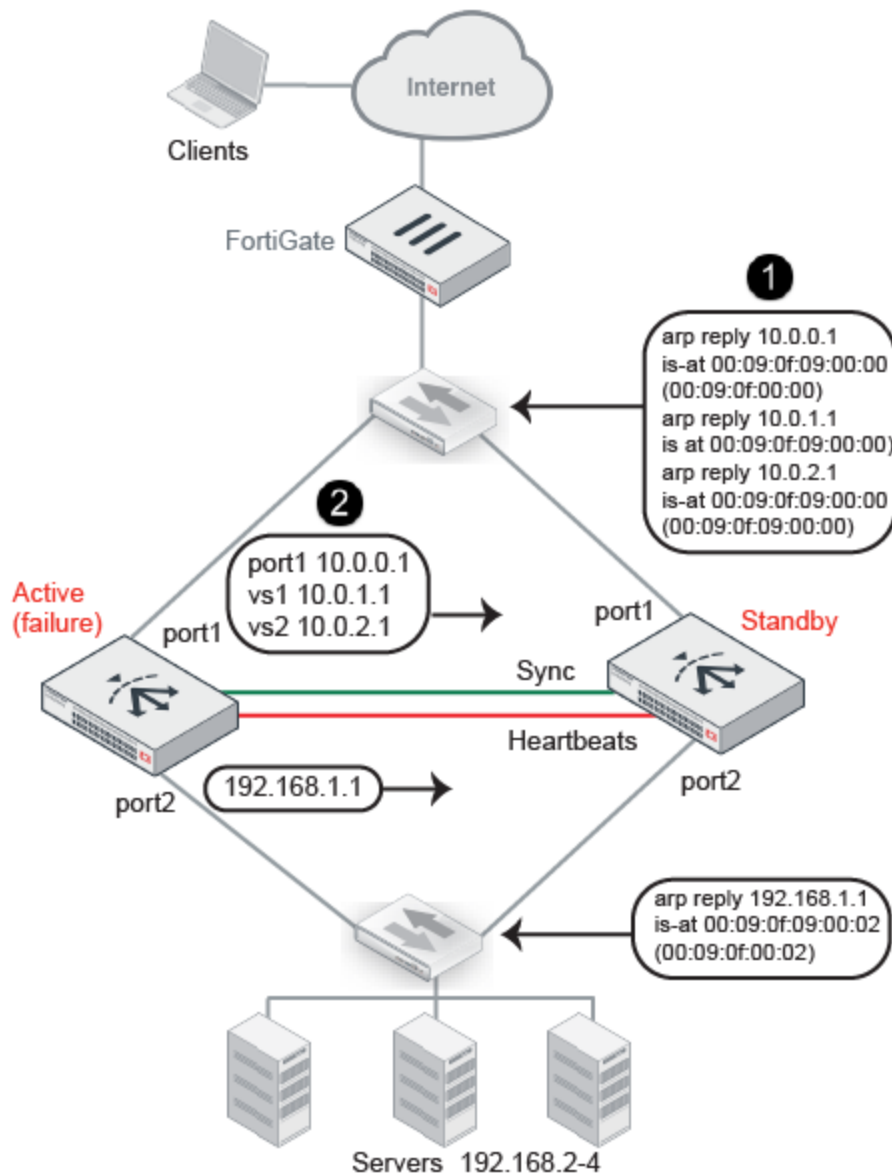
This topic includes the following information:

- [Overview](#)
- [Basic steps](#)
- [Best practice tips](#)

Overview

In an active-passive cluster, one node is the active appliance; it processes traffic. The other node is passive; it is ready to assume the role of the active appliance if the primary node is unavailable.

You configure the system to send heartbeat packets between the pair to monitor availability. The system continually polls the activity of the heartbeat packets. If the active appliance becomes unresponsive, failover occurs: the standby becomes active. [Figure 73](#) illustrates the process: (1) the standby node sends gratuitous ARP to notify adjacent routers to direct traffic for the virtual MAC addresses (vMAC) to its network interfaces; (2) It takes the IP addresses of the unresponsive node.

Figure 73: An active-passive cluster at failover—IP address transfer to the new active member

When the former active appliance comes back online, it might or might not assume its former active role. The system selects the active member based on the following criteria:

- Link health (if monitor ports links are down, the node is considered down)
- Remote IP monitor health check results
- Override setting (prefers priority to uptime)
- Most available ports
- Highest uptime value
- Lowest device priority number (1 has greater priority than 2)
- Highest-sorting serial number—Serial numbers are sorted by comparing each character from left to right, where 9 and z are the greatest values. The system gives preference to higher values over lower values.

Basic steps

To deploy an active-passive cluster:

1. License all FortiADC appliances in the HA cluster, and register them, including FortiGuard services, with the Fortinet Customer Service & Support website:
<https://support.fortinet.com/>
2. Physically link the FortiADC appliances that make up the HA cluster.
You must link at least one of their ports (for example, port4 to port4) for heartbeat and synchronization traffic between members of the cluster. You can do either of the following:
 - Connect the two appliances directly with a crossover cable.
 - Link the appliances through a switch. If connected through a switch, the heartbeat interfaces must be reachable by Layer 2 multicast.
3. Configure the secondary node:
 - a. Log into the secondary appliance as the **admin** user.
 - b. Complete the HA settings as described in [Configuring HA settings](#).
Important: Set the Device Priority to a higher number than the preferred primary node; for example, set it to 2.
4. Configure the primary node:
 - a. Log into the primary appliance as the **admin** user.
 - b. Complete the configuration for all features, as well as the HA configuration.
Important: Set the Device Priority to a lower number than the secondary node; for example, set it to 1.

Note: After you have saved the HA configuration changes, cluster members join or rejoin the cluster. After you have saved configuration changes on the primary node, it automatically pushes its configuration to the secondary node.

Best practice tips

The following tips are best practices:

- Be careful to maintain the heartbeat link(s). If the heartbeat is accidentally interrupted, such as when a network cable is temporarily disconnected, the other nodes assume that the primary node has failed. In an active-passive deployment, failover occurs. If no failure has actually occurred, both nodes can be operating as the active node simultaneously.
- If you link HA appliances through switches, to improve fault tolerance and reliability, link the ports through two *separate* switches. Also, do *not* connect these switches to your overall network, which could introduce a potential attack point, and could also allow network load to cause latency in the heartbeat, which could cause an unintentional failover.

Deploying an active-active cluster

This topic includes the following information:

- [Configuration overview](#)
- [Basic steps](#)
- [Expected behavior](#)
- [Best practice tips](#)

Configuration overview

Figure 74 shows an example of an active-active cluster. In an active-active cluster, traffic from the upstream router can be load-balanced among up to eight member nodes.

Load balancing depends on the equal cost multipath (ECMP) configuration on adjacent routers. The routers on either side of the cluster must be configured to use ECMP to distribute traffic to the FortiADC cluster nodes. In the example, assume that the FortiADC configuration includes virtual servers belonging to subnet 10.61.0.0/24. On Router A, you configure equal cost routes as follows:

```
destination: 10.61.0.0/24 gateway: 10.61.51.1
destination: 10.61.0.0/24 gateway: 10.61.51.2
destination: 10.61.0.0/24 gateway: 10.61.51.3
```

Likewise, on Router B, you configure equal cost routes for server-to-client traffic:

```
destination: 0.0.0.0/0 gateway: 10.65.51.1
destination: 0.0.0.0/0 gateway: 10.65.51.2
destination: 0.0.0.0/0 gateway: 10.65.51.3
```

Active-active clusters also support *failover*. The primary node is the backup node for each of the other nodes in the cluster. If a member node fails, the primary node takes its IP address and sends gratuitous ARP to adjacent routers to direct traffic for that virtual MAC address (vMAC) to its own network interfaces.

The FortiADC configuration involves the following components:

- Primary node system and feature configuration
- Interface configuration (HA node IP list)
- HA configuration

In an active-active cluster, one of the nodes is selected as the *primary node*, and the others are *member nodes*. In this example, node1 is the primary node and node2 and node3 are member nodes. When the cluster is formed, the configuration for node1 is pushed to node2 and node3.

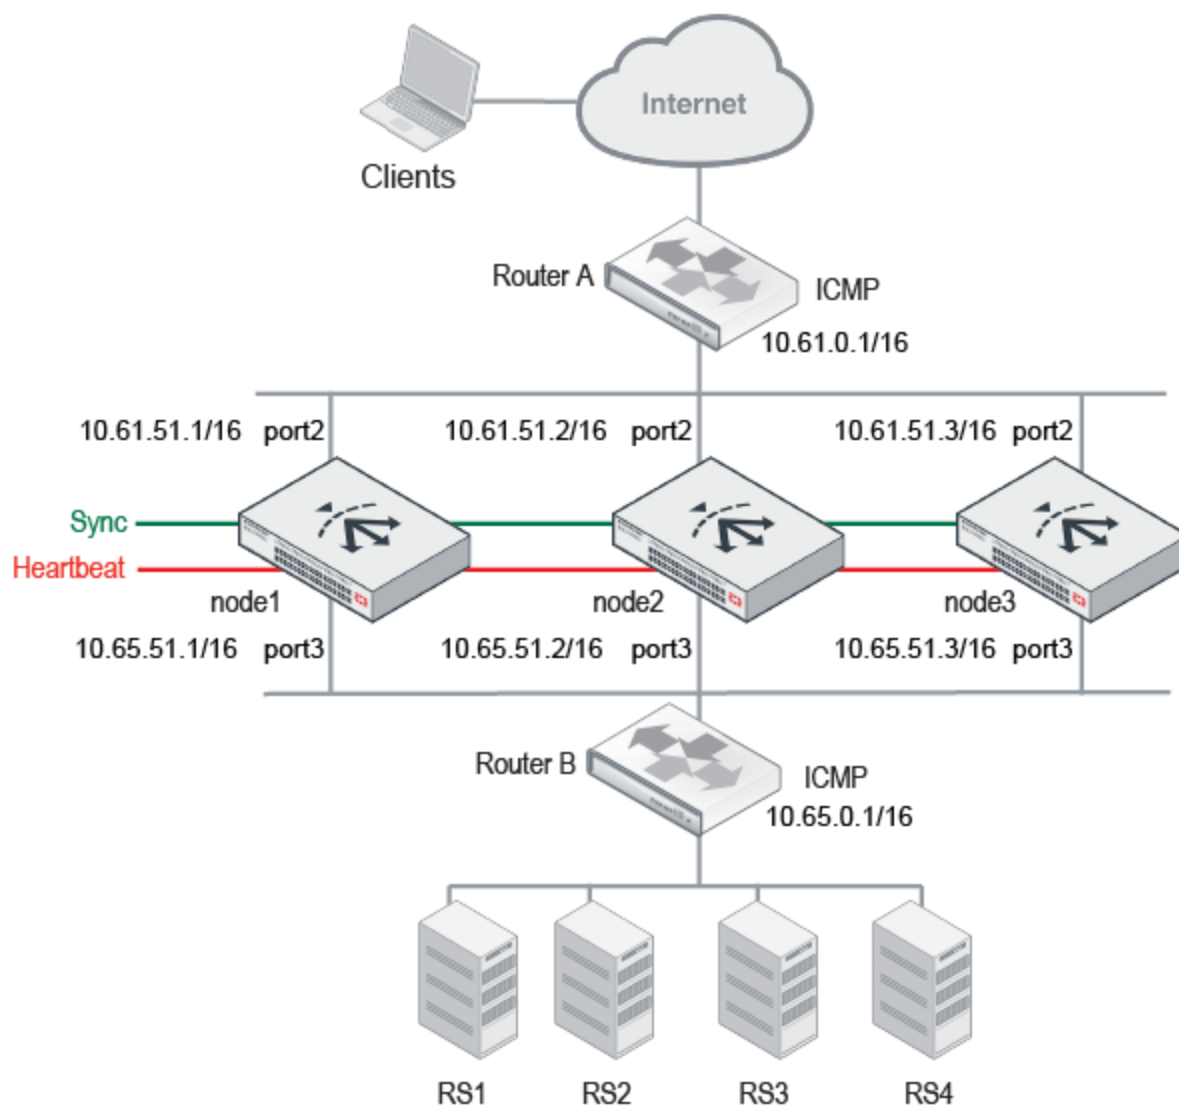
When you configure the network interfaces for nodes in an active-active cluster, in addition to the interface primary IP address, you configure an HA node IP list that specifies special HA IP addresses of each node in the cluster. The HA node IP list for port2 in the example has the following values:

```
10.61.51.1/16 node1
10.61.51.2/16 node2
10.61.51.3/16 node3
```

Likewise, the HA node IP list for port3 has the following values:

```
10.65.51.1/16 node1
10.65.51.2/16 node2
10.65.51.3/16 node3
```

Finally, you log into each node when it is in standalone mode to configure its HA settings. When you are ready to form the cluster, change the setting to HA active-active. The system state changes when a node joins a cluster.

Figure 74: HA active-active deployment

Note: The example shows routers on both sides of the FortiADC cluster. Your deployment might not have a router between the FortiADC cluster and the real server pool. In this case, if your real servers support load balancing methods like ECMP, the expected behavior is the same as what is described here. If not, it is expected that the real servers route reply traffic to the cluster node that sent them the client traffic.

Basic steps

To deploy an active-active cluster:

1. License all FortiADC appliances in the HA cluster, and register them, including FortiGuard services, with the Fortinet Customer Service & Support website: <https://support.fortinet.com/>.
2. Physically link the FortiADC appliances that make up the HA cluster.
You must link at least one of their ports (for example, port4 to port4) for heartbeat and synchronization traffic between members of the cluster. You can do either of the following:

- If only two nodes, connect the two appliances directly with a crossover cable.
 - If more than two nodes, link the appliances through a switch. If connected through a switch, the interfaces must be reachable by Layer 2 multicast.
3. Configure member nodes:
 - a. Log into the member nodes as the **admin** user.
 - b. Complete the HA configuration as described in [Configuring HA settings](#).
Important: Set the Device Priority to a higher number than the preferred primary node; for example, set it to 2.
 4. Configure the preferred primary node:
 - a. Log into the primary node as the **admin** user.
 - b. Configure network interfaces so that each traffic interface has an HA node IP address list in addition to its physical port IP address. See [Configuring network interfaces](#).
When HA is set to standalone, the system uses the physical port IP address. When HA is set to active-active, the system uses the HA node IP address.
 - c. Complete the configuration for all features, as well as the HA configuration.
Important: Set Device Priority to a lower number than the member nodes; for example, set it to 1.

Note: After you have saved the HA configuration changes, cluster members join or rejoin the cluster. After you have saved configuration changes on the primary node, it automatically pushes its configuration to the member nodes.

Expected behavior

In active-active deployments, be sure to enable data synchronization. In particular, enable the following settings:

- Layer 4 Connection Synchronization—Synchronizes TCP connection state data.
- Layer 4 Session Synchronization—Synchronizes the source IP address table used for persistence to backend servers.
- Layer 7 Session Synchronization—Synchronizes the source IP address table used for persistence to backend servers.

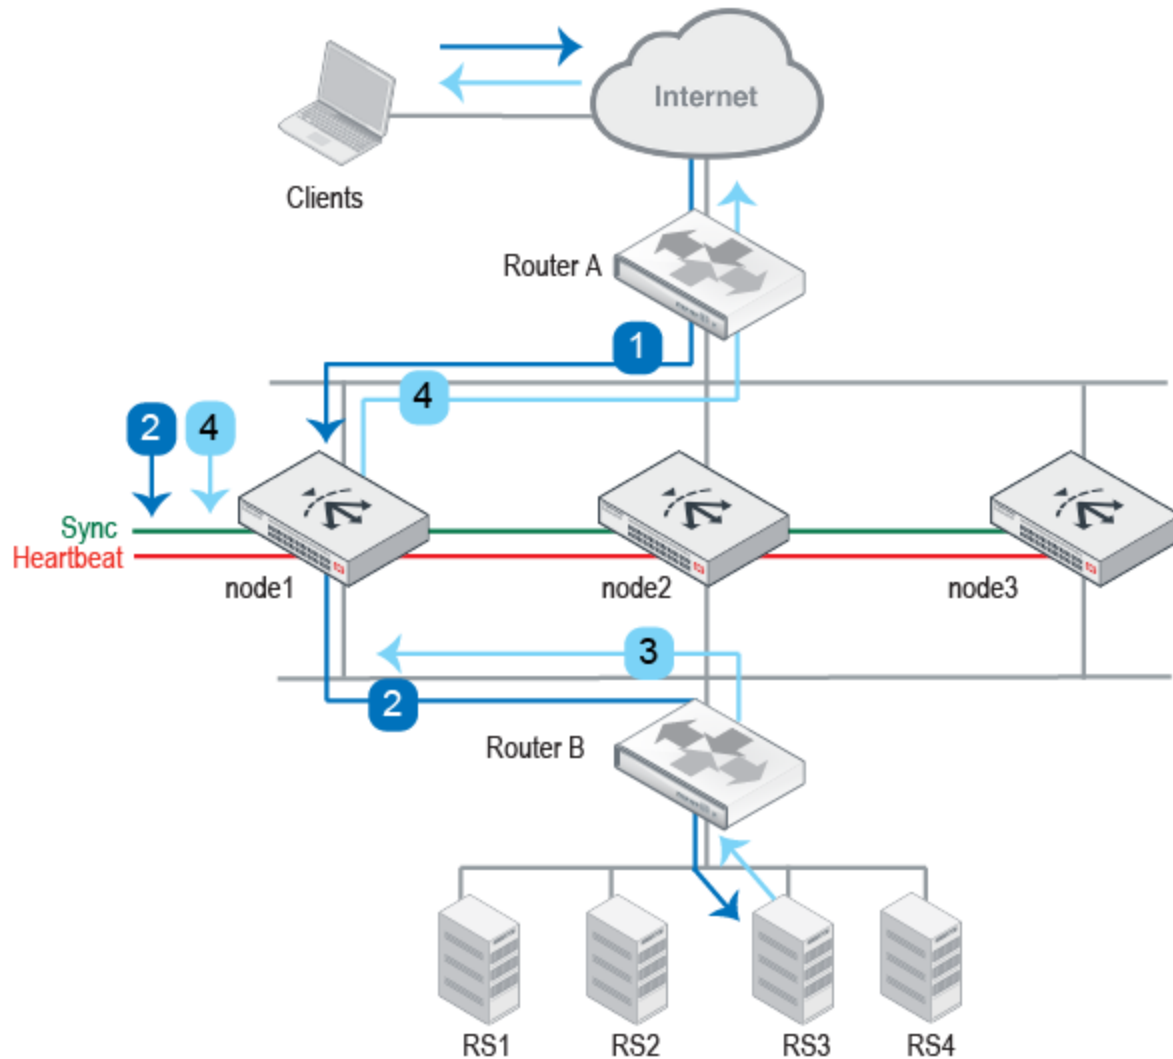
The sections that follow describe how the cluster uses synchronized data.

Traffic to TCP virtual servers

When Layer 4 synchronization is enabled, the cluster nodes share TCP connection state and Layer 4 source IP address data for traffic to Layer 4 virtual servers (and Layer 2 TCP and Turbo HTTP virtual servers, which are packet-based). The node that receives the first SYN packet forwards the traffic to the real server, and, at the same time, multicasts the session data to the other nodes in the cluster.

Figure 75 illustrates the sequence of the traffic flow when client-to-server and server-to-client session traffic are routed through the same node.

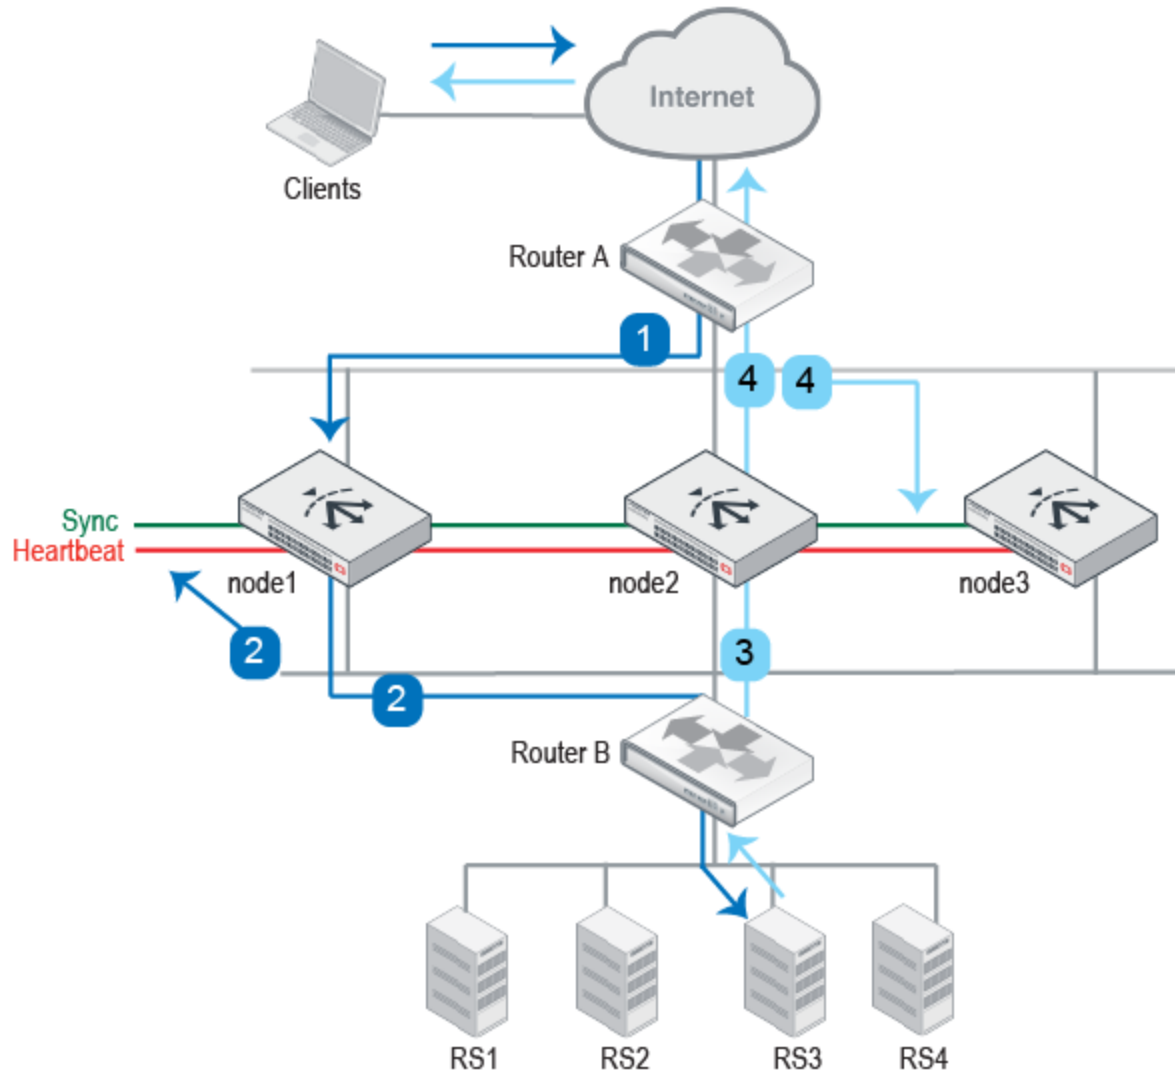
Figure 75: TCP traffic flow when ECMP results in forwarding through same node



1. Router A uses ECMP to select a cluster node to which to forward a client connection request—in this case, node1.
2. The cluster node forwards the traffic to a real server and multicasts the session data to the cluster via the data port.
3. Router B uses ECMP to select a cluster node to which to forward the server response traffic—also node1.
4. The cluster node forwards the traffic to the client and multicasts the session data to the cluster.

Figure 76 illustrates the sequence of the traffic flow when client-to-server and server-to-client session traffic are routed through different nodes and synchronization has occurred before the second node receives the response traffic.

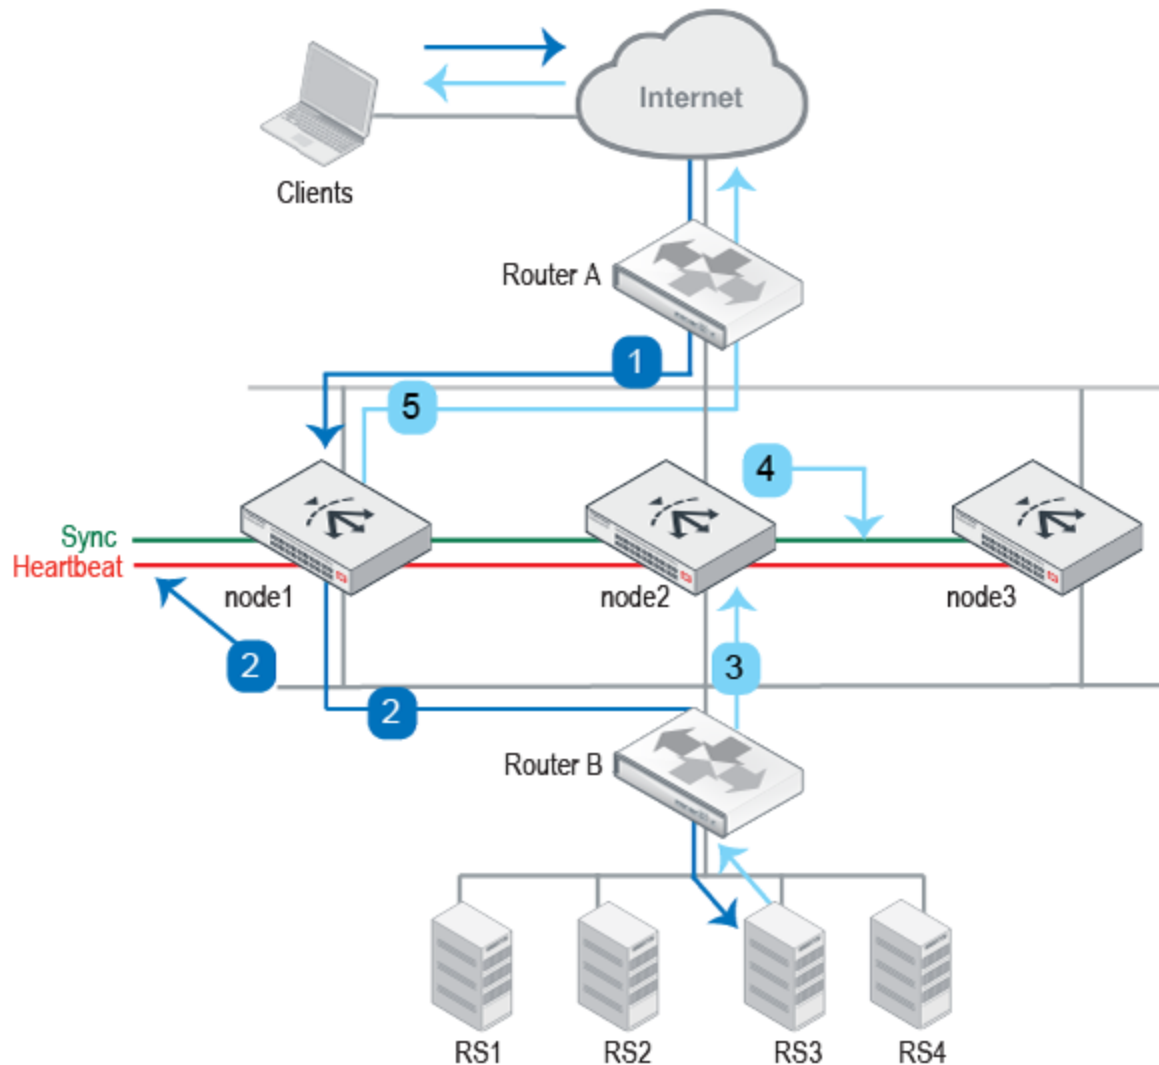
Figure 76: TCP traffic flow when synchronization has occurred



1. Router A uses ECMP to select a cluster node to which to forward a client connection request—in this case, node1.
2. The cluster node forwards the traffic to a real server and multicasts the session data to the cluster via the data port.
3. Router B uses ECMP to select a cluster node to which to forward the server response traffic. In this case, it selects node2.
4. If the session has already been synchronized between node1 and node2, node2 forwards the traffic to the client and multicasts the session data to the cluster.

Figure 77 illustrates the sequence of the traffic flow when client-to-server and server-to-client session traffic are routed through different nodes and synchronization has not yet occurred when the second node receives the response traffic.

Figure 77: TCP traffic flow when synchronization has not yet occurred



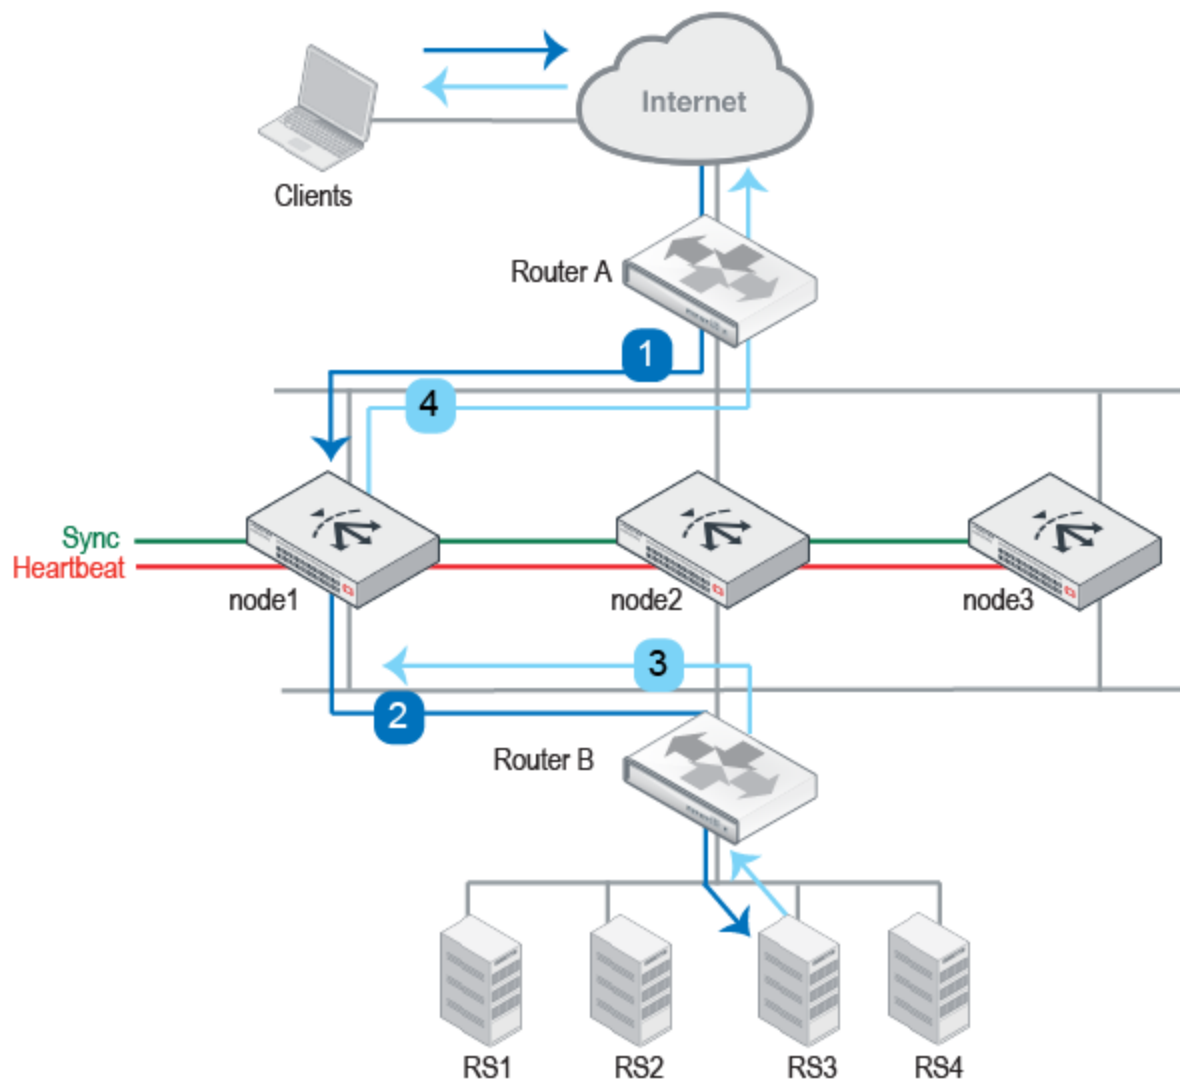
1. Router A uses ECMP to select a cluster node to which to forward a client connection request—in this case, node1.
2. The cluster node forwards the traffic to a real server and multicasts the session data.
3. Router B uses ECMP to select a cluster node to which to forward the server response traffic. In this case, it selects node2.
4. Because the session has not yet been synchronized between node1 and node2, node2 multicasts the traffic to the cluster.
5. When node1 receives traffic from node2, it forwards the traffic to the client and multicasts the session data.

Traffic to HTTP virtual servers

When Layer 7 synchronization is enabled, the cluster nodes share source IP data for traffic to HTTP virtual servers differently when the virtual server profile Source option is enabled. When the Source option is enabled, the traffic FortiADC forwards to the real server has the client source IP address; when disabled, it has the FortiADC HA cluster node IP address.

Figure 78 illustrates the sequence of the traffic flow when the Source option is not enabled.

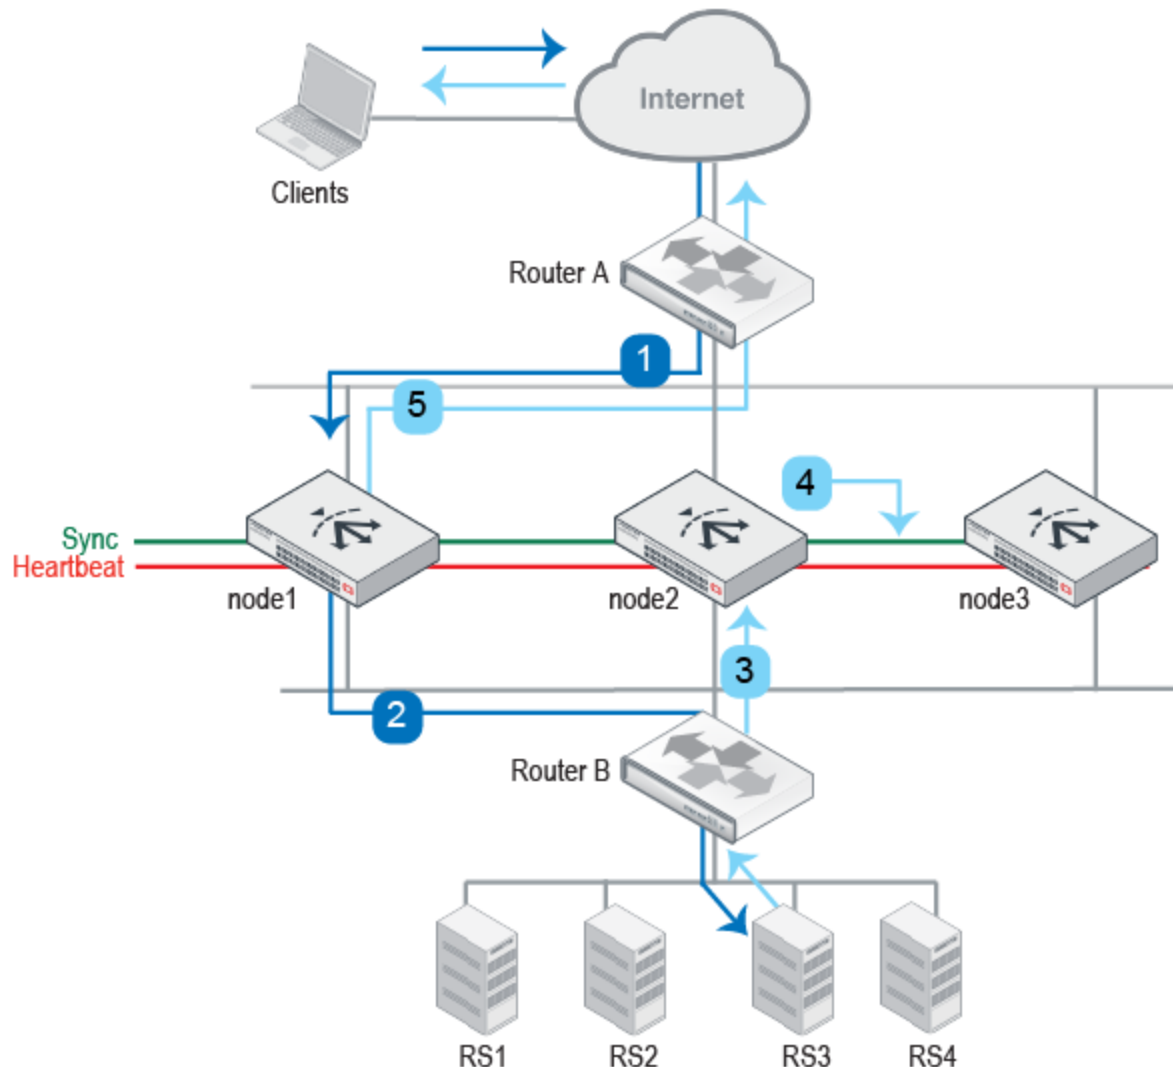
Figure 78: HTTP traffic flow when the Source profile option is not enabled



1. Router A uses ECMP to select a cluster node to which to forward a client connection request—in this case, node1.
2. The cluster node forwards the traffic to a real server. Because the Source option was not enabled, the source IP address in the FortiADC-to-real-server traffic is the node1 HA cluster node IP address, and this becomes the destination IP address for the response traffic.
3. Router B does not use ECMP; instead, it forwards the traffic to the node1 HA cluster IP address.
4. The cluster node finds the real client IP address in its session table and forwards the traffic to the client.

Figure 79 illustrates the sequence of the traffic flow when the Source option is enabled.

Figure 79: HTTP traffic flow when the Source profile option is enabled



1. Router A uses ECMP to select a cluster node to which to forward a client connection request—in this case, node1.
2. The cluster node forwards the traffic to a real server. Because the Source option is enabled, the source IP address in the FortiADC-to-real-server traffic is the true client IP address, and this becomes the destination IP address for the server-to-client traffic.
3. Router B uses ECMP and might forward the traffic to any node in the cluster. In this example, it forwards the traffic to node2.
4. Because the server-to-client response was not expected by node2, it multicasts the traffic to the cluster.
5. When node1 receives the server-to-client response data from node2, it forwards the response to the client.

Note: In an active-active deployment, the virtual server profile Source option adds latency to the transaction. To reduce latency, use an alternative to the Source option, such as the X-Forwarded-For option, if you have a requirement that the original client IP be logged by the real server.

FTP traffic and traffic processed by firewall rules

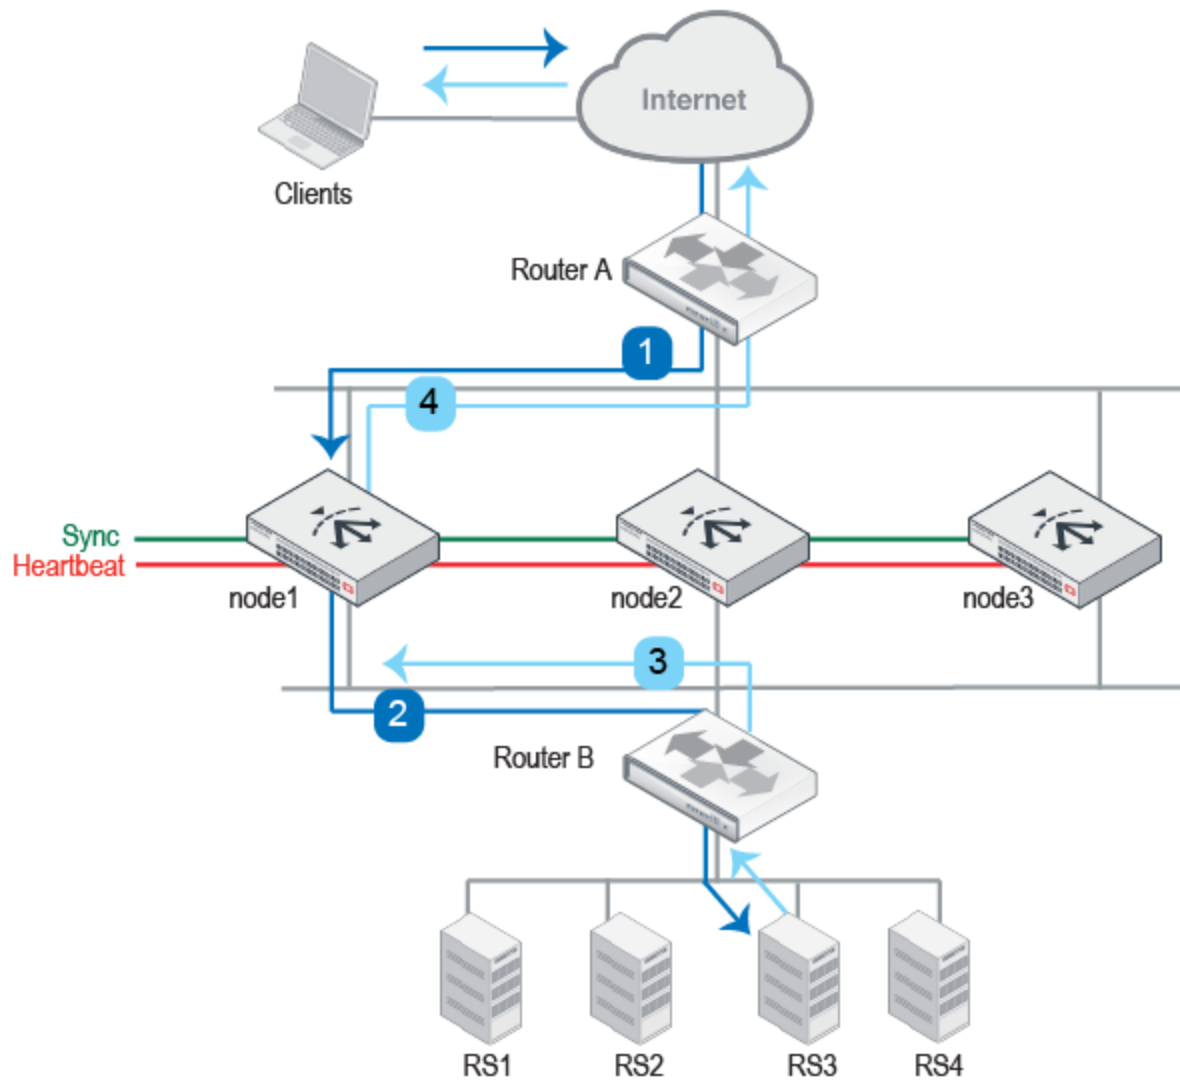
In an active-active deployment, FTP traffic and firewall traffic are always forwarded through the primary node only.

FTP has both a control connection and a data connection associated with client-server communication. The two “channels” make it difficult to support asymmetric routes in an active-active cluster.

In addition, traffic processed by the stateful firewall rules is also not load-balanced.

Figure 80 illustrates the sequence of the traffic flow when ECMP results in traffic being forwarded through the primary node.

Figure 80: FTP or firewall traffic flow when ECMP selects the primary node

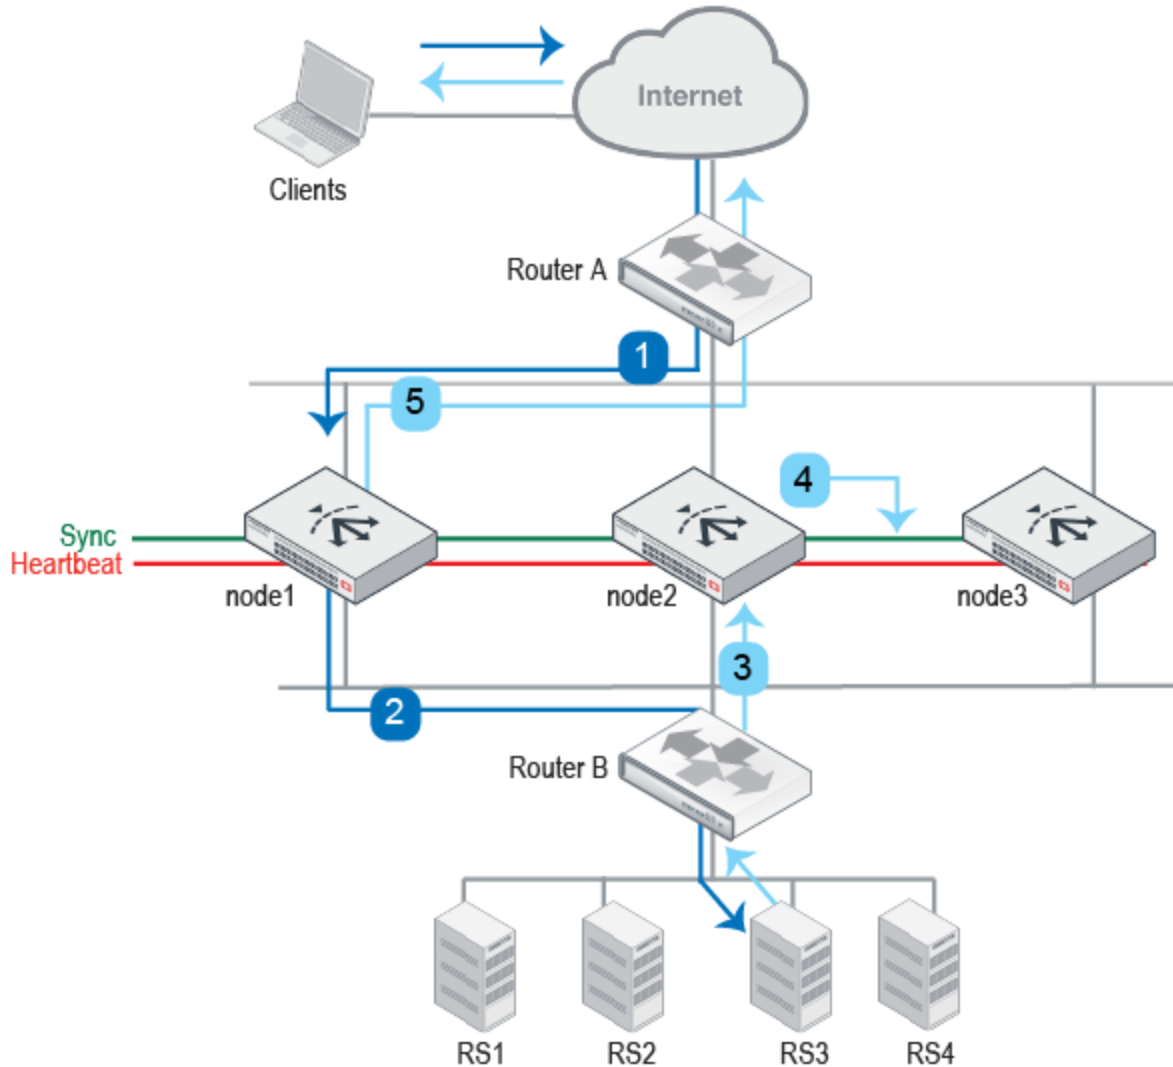


1. Router A uses ECMP to select a cluster node to which to forward a client connection request. In this case, it selects the primary node, node1.
2. The primary node forwards the traffic to a real server.

3. Router B uses ECMP to select a cluster node to which to forward the server response traffic—also node1.
4. The primary node forwards the traffic to the client.

Figure 81 illustrates the sequence of the traffic flow when ECMP results in an asymmetric route.

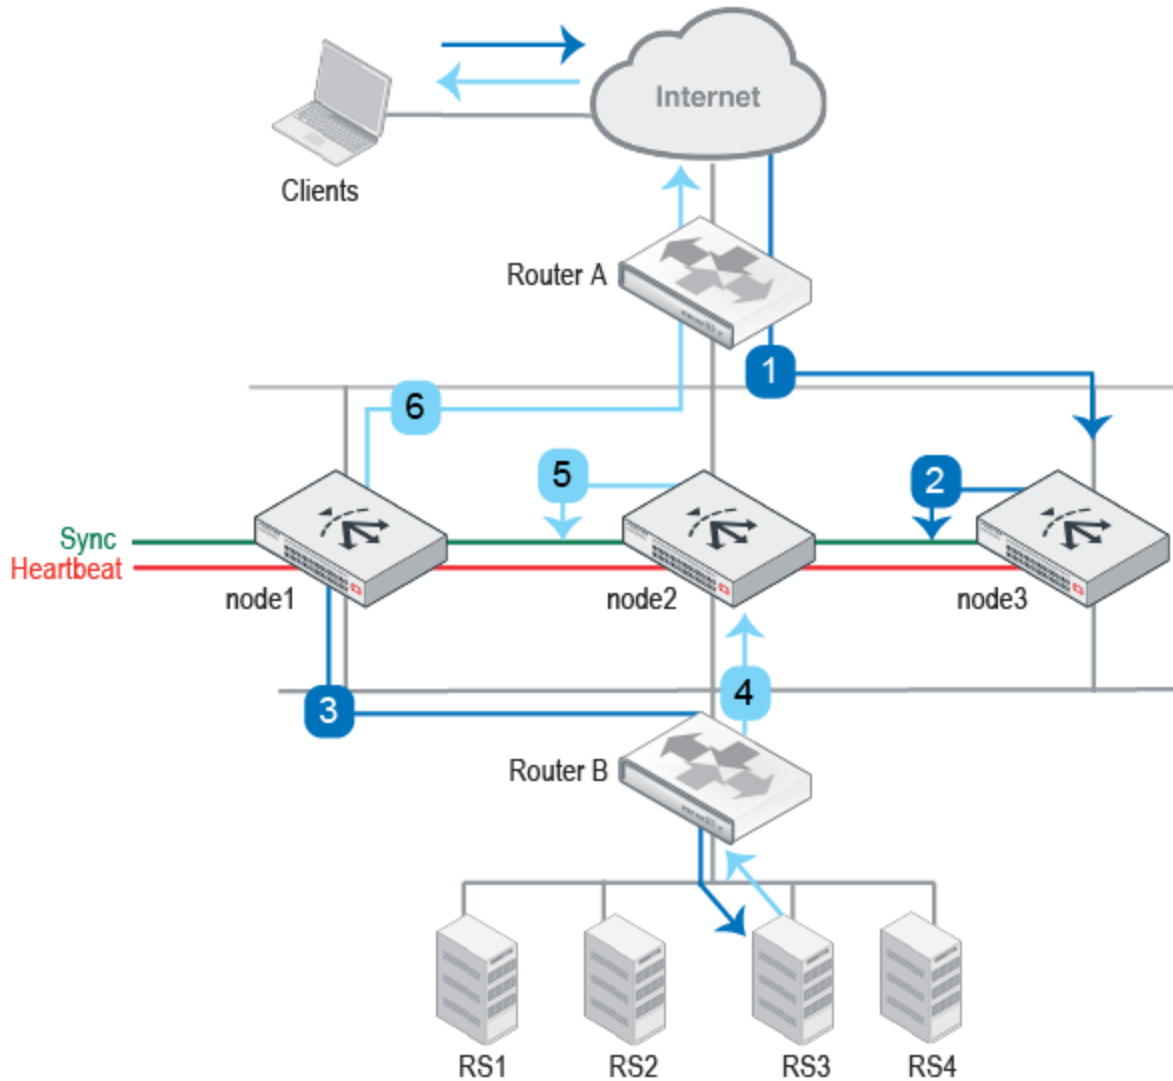
Figure 81: FTP or firewall traffic flow when ECMP results in an asymmetric route



1. Router A uses ECMP to select a cluster node to which to forward a client connection request. In this case, it selects the primary node, node1.
2. The cluster node forwards the traffic to a real server.
3. Router B uses ECMP to select a cluster node to which to forward the server response traffic—in this case, node2.
4. Because the server-to-client response was not expected by node2, it forwards traffic to the cluster.
5. When the primary node receives traffic from node2, it forwards it to the client.

Figure 82 illustrates the sequence of the traffic flow when ECMP results in client-to-server traffic sent to a non-primary node.

Figure 82: FTP or firewall traffic flow when ECMP results in traffic sent to a non-primary node



1. Router A uses ECMP to select a cluster node to which to forward a client connection request to a real server destination IP address. In this case, it selects a member node, node3.
2. Firewall traffic is forwarded by the primary node only, so node3 multicasts the session data to the cluster.
3. The primary node forwards the traffic to a real server.
4. Router B uses ECMP to select a cluster node to which to forward the server response traffic—in this case, node2.
5. Because the server-to-client response was not expected by node2, it forwards traffic to the cluster.
6. When the primary node receives traffic from node2, it forwards it to the client.

Best practice tips

The following tips are best practices:

- Be careful to maintain the heartbeat link(s). If the heartbeat is accidentally interrupted, such as when a network cable is temporarily disconnected, the other nodes assume that the primary node has failed. In an active-active deployment, a new primary node is elected among member nodes. If no failure has actually occurred, both nodes can be operating as primary nodes simultaneously.
- If you link HA appliances through switches, to improve fault tolerance and reliability, link the ports through two *separate* switches. Also, do *not* connect these switches to your overall network, which could introduce a potential attack point, and could also allow network load to cause latency in the heartbeat, which could cause an unintentional failover.

Advantages of HA Active-Active-VRRP

Compared with HA Active-Passive or Active-Active clusters, an HA Active-Active-VRRP cluster offers the following advantages:

- The HA Active-Active mode is an device-based HA mode, in which the HA fail over will switch over the whole failed device even in cases where only one monitor port fails.
- In FortiADC HA Active-Active-VRRP mode, you can manually assign a virtual server to a traffic group, enabling you to do traffic load design based on virtual servers.
- In HA Active-Active-VRRP mode, FortiADC only synchronizes the session table/persistence table to the next available device in the same traffic group using the “failover-order “ command. In cases where you have more than two devices in the cluster, this synchronization mechanism can turn out to be more efficient than HA Active-Passive or Active-Active mode because the session/persistence table will be synced to the whole HA group. In this sense, FortiADC actually supports the N+M hot-backup function.
- HA Active-Active mode must work together with an external router with the ECMP route configured to distribute traffic to different Active-Active nodes; HA Active-Active-VRRP mode does not need this external router to do ECMP traffic distribution — Both sides can simply point their respective gateway to the VRRP floating IP.
- In HA Active-Active-VRRP mode, different devices in the same traffic group have the same HA status. Once you have pointed both the client and the server side gateways to the floating IP in the same traffic, the incoming/outgoing traffic will going to the same device. As a result, HA Active-Active-VRRP mode doesn't need to multicast the traffic itself to the HA group, which should offer the best network performance and efficiency.
- In HA Active-Active mode, the AA-Master will take over all AA-NotWorking nodes' traffic. If multiple AA devices have failed, the AA-Master will have to process much more traffic than the AA-Slave nodes, which may exhibit some unexpected behavior under abnormal high traffic stress.
- In terms of sync session, you are unable to access the real server's IP address from the client directly in HA Active-Active mode, but you don't have this limitation in HA Active-Active-VRRP mode.

Deploying an active-active-VRRP cluster

This topic includes the following information:

- [Configuration overview](#)
- [Basic steps](#)

- [Best practice tips](#)

Configuration overview

The Virtual Router Redundancy Protocol (VRRP) is designed to eliminate the single point of failure inherent in the static default routed environment. VRRP specifies an election protocol that dynamically assigns responsibility for a virtual router to one of the VRRP routers on a LAN. The VRRP router controlling the IP address(es) associated with a virtual router is called the Master, and forwards packets sent to these IP addresses. The election process provides dynamic fail-over in the forwarding responsibility should the Master become unavailable. Any of the virtual router's IP addresses on a LAN can then be used as the default first hop router by end-hosts. The advantage of VRRP is a higher availability default path without requiring configuration of dynamic routing or router discovery protocols on every end-host.

A virtual router is defined by its virtual router identifier (VRID) and a set of IP addresses. A VRRP router may associate a virtual router with its real address on an interface, and may also be configured with additional virtual router mappings and priority that the virtual router can back up. The mapping between VRID and addresses must be coordinated among all VRRP routers on a LAN.


FortiADC only adopts the VRRP concept, but not the exact VRRP protocol itself. For this reason, its HA Active-Active VRRP mode can only be called a VRRP-like HA mode.

VRRP configurations can be used as a high availability (HA) solution to ensure that your network maintains connectivity with the Internet (or with other networks) even if the default router for your network fails. Using VRRP, you can assign VRRP routers as master or backup routers. The master router processes traffic, while the backup routers monitor the master router and start forwarding traffic the moment the master router fails.

VRRP is described in RFC 3768.

FortiADC units can function as master or backup Virtual Router Redundancy Protocol (VRRP) routers and can be quickly and easily integrated into a network that has already deployed VRRP. In a VRRP configuration, when a FortiADC unit operating as the master unit fails, a backup unit automatically takes its place and continues processing network traffic. In such a situation, all traffic to the failed unit transparently fails over to the backup unit that takes over the role of the failed master FortiADC unit. When the failed FortiADC unit is restored, it will once again take over processing traffic for the network.

Figure 83: An active-active-VRRP cluster configuration using two FortiADC units



In an active-active-VRRP cluster, one of the nodes is selected as the primary node of a traffic group, and the rest of the nodes are member nodes of the traffic group. Traffic from the upstream can be load-balanced among up to eight member nodes. Active-active-VRRP clusters also support failover. If the primary node fails, the traffic group work on this node will fail over to one of the backup nodes which will send gratuitous ARP to adjacent devices to redirect traffic for its own MAC address to all network interfaces within the traffic group.

The FortiADC VRRP configuration involves the following:

- Traffic group and their features
- Interface and virtual server (pertinent floating IP and traffic group)
- HA

Note:It is important to note that FortiADC only supports VRRP configuration between two or more FortiADC units. It can NOT be integrated into a VRRP group formed with any third-party VRRP devices.

Basic steps

To deploy an active-active-VRRP cluster:

1. Configure the HA active-active-VRRP cluster.
<https://support.fortinet.com/>

For example:

```
config system ha
set mode active-active-vrrp
set hbdev port2
```

```

        set group-id 14
        set local-node-id 1
    end

```

2. Configure the traffic group.

Configure the traffic group and set its parameters. The failover sequence must be configured according to the order of node IDs. This means that if a node is dead, the next node in queue will take over handling the traffic. If you want to decide when a node should retake the traffic over from power-down to start-up, you can enable the preempt.

If only two nodes, connect the two appliances directly with a crossover cable.

If more than two nodes, link the appliances through a switch. If connected through a switch, the interfaces must be reachable by Layer 2 multicast.

```

config system traffic-group
    edit "traffic-group-1"
        set failover-order 1 2
    next
end

```

3. Configure applications and relate them with the traffic group

Relate applications with the traffic group in the virtual server configuration and interface + IP configuration. If no traffic group is related, the “default” traffic group will be used.

For example (Relate a virtual server to a traffic group):

```

config load-balance virtual-server
    edit "vs1"
        set packet-forwarding-method FullNAT
        set interface port1
        set ip 10.128.3.4
        set load-balance-profile LB_PROF_HTTP
        set load-balance-method LB_METHOD_DEST_IP_HASH
        set load-balance-pool rs1
        set ippool-list vs1-pool vs1-pool-1
        set traffic-group traffic-group-1
    next

```

For example (Relate an interface and IP address with a traffic group):

```

edit "port1"
    set vdom root
    set ip 10.128.3.1/16
    set allowaccess https ping ssh snmp http telnet
    set traffic-group traffic-group-1
    set floating enable
    set floating-ip 10.128.3.3
next
end

```

Best practice tips

The following tips are best practices:

Note: After you have saved the HA configuration changes, cluster members join or rejoin the cluster. After you have saved configuration changes on the primary node, it automatically pushes its configuration to the

member nodes.

Chapter 15: Virtual Domains

This chapter includes the following topics:

- [Virtual domain basics](#)
- [Enabling the virtual domain feature](#)
- [Creating virtual domains](#)
- [Assigning network interfaces and admin users to VDOMs](#)
- [Virtual domain policies](#)
- [Disabling virtual domains](#)

Virtual domain basics

A virtual domain (VDOM) is a complete FortiADC instance that runs on the FortiADC platform. The VDOM feature supports multitenant deployments. To do this, you create a virtual domain configuration object that contains all of the system and feature configuration options of a full FortiADC instance, and you provision an administrator account with privileges to access and manage only that VDOM.

Note: The super user **admin** can access all VDOMs that have been created on the system, but the administrator accounts that are provisioned for a VDOM can access only that particular VDOM.

To use the VDOM feature, complete the following steps:

1. Enable the virtual domain feature.
2. Create a virtual domain configuration object.
3. Assign network interfaces and administrators to the virtual domain.

Enabling the virtual domain feature

You can use the web UI to enable the virtual domain feature. By default, the virtual domain feature is not enabled, and the GUI for virtual domain configuration is hidden.

Before you begin:

- You must have super user permission (user **admin**) to enable the virtual domain feature.

To enable the virtual domain feature:


1. Go to System > Settings.
The configuration page displays the Basic tab.
2. Enable **Virtual Domain**.
3. Save the configuration.

Note: You can also enable the virtual domain feature from the System Information panel under Dashboard > Status.

Figure 84 shows the landing page after the **admin** administrator logs into the system when the virtual domain feature is enabled. From here, the **admin** administrator can create virtual domains, assign network interfaces to virtual domains, create admin users for virtual domains, and navigate to the system and feature configuration pages for the virtual domains, including the root (default) domain.

When a user with a delegated administrator account logs in, the landing page is the standard landing page. These users cannot perform the tasks related to virtual domain administration that the **admin** administrator performs.

Figure 84: Super admin login with virtual domain



Creating virtual domains

By default, FortiADC has a predefined virtual domain named root that you cannot delete or modify. The **admin** user can add, delete, enable, and disable virtual domains.

Before you begin:

- You must have super user permission (user **admin**) to create virtual domains.
- You must have super user permission (user **admin**) to assign network interfaces to virtual domains.

To create a virtual domain:

1. Go to Virtual Domain.
2. Click **Add**, enter a unique name for the virtual domain.
3. Save the configuration.

Assigning network interfaces and admin users to VDOMs

By default, all network interfaces are assigned to the root virtual domain. After you have created the virtual domain, you can assign network interfaces to it.

To assign a network interface to a virtual domain:

1. Go to Networking > Interface.
2. Double-click an interface configuration or click **Add** to create one.
3. Configure interface settings and select the virtual domain.
4. Save the configuration.

When virtual domain administrators log into the FortiADC system, they only see configuration settings and data for the virtual domain that you assigned them to. They do not see the Virtual Domains menu in the navigation pane.


To create an administrator for a virtual domain:

1. Go to System > Administrator.
2. Click **Add** to create an administrator.
3. Configure administrator settings and select the virtual domain.
4. Save the configuration.

Virtual domain policies

FortiADC allows you to create and impose custom policies or restrictions on each virtual domain you have added. For each virtual domain, you can configure the maximum range for its Dynamic Resources and Static Resources. Dynamic Resources are related to a virtual domain's performance, while Static Resources are related to its configuration. The Vdom configuration dialog ([Figure 85](#)) also shows a virtual domain's current configuration and workload settings, which serve as good reference points for you to fine-tune the virtual domain.

Figure 85: Vdom configuration

 Vdom

Name

Dynamic Resources

	Current	Max (Range:0-1000000)
L4 CPS	0.0	<input type="text" value="0"/>
L7 CPS	0.0	<input type="text" value="0"/>
L7 RPS	0.0	<input type="text" value="0"/>
SSL CPS	0.0	<input type="text" value="0"/>
SSL Throughput	0.0	<input type="text" value="0"/>
Concurrent Session	0.0	<input type="text" value="0"/>

Static Resources

	Current	Max (Range:0-1024)
Virtual Server	0	<input type="text" value="0"/>
Real Server	0	<input type="text" value="0"/>
Health Check	4	<input type="text" value="0"/>
Source Pool	0	<input type="text" value="0"/>
Error-Page	0	<input type="text" value="0"/>
Local User	0	<input type="text" value="0"/>
User Group	0	<input type="text" value="0"/>

Disabling virtual domains

To disable the virtual domain feature:

1. Assign all network interfaces and administrators to the **root** virtual domain.
2. Delete all virtual domains.
3. Clear the **Virtual Domain** option.

Chapter 16: SSL Transactions

This chapter covers the following topics:

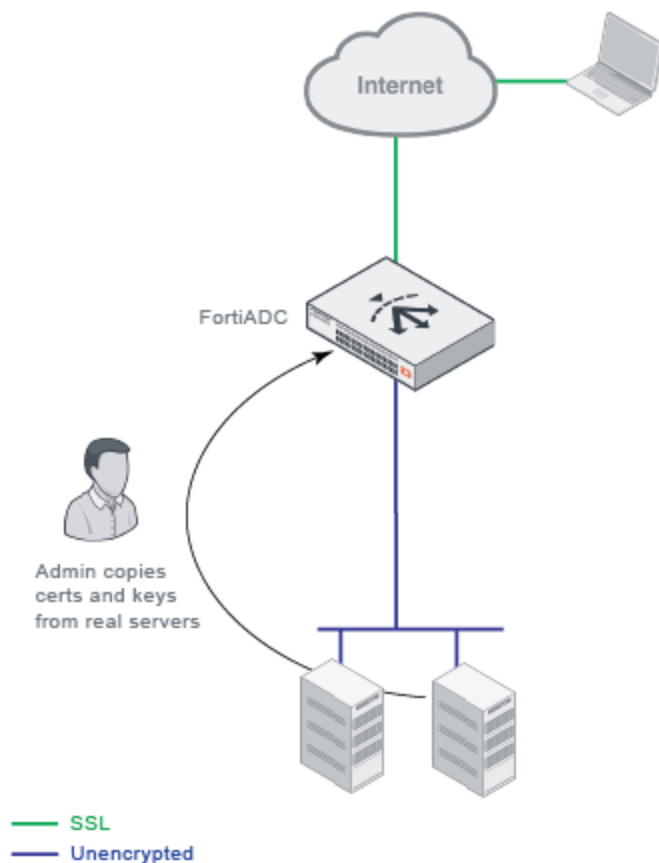
- [SSL offloading](#)
- [SSL decryption by forward proxy](#)
- [Profile configurations](#)
- [Certificate guidelines](#)
- [SSL/TLS versions and cipher suites](#)
- [Exceptions list](#)
- [SSL/HTTP mirror traffic](#)

SSL offloading

You can use FortiADC in a Layer-7 load-balancing topology to offload SSL decryption from the real server farm. In such a deployment, the FortiADC unit uses a copy of the real server certificate and its private key to negotiate the SSL connection. It acts as an SSL proxy for the real servers, using the certificates and their private keys to:

- authenticate itself to clients
- decrypt requests
- encrypt responses

When session data has been decrypted, you can use the FortiADC content rewriting, content routing, and web application firewall features.

Figure 86: SSL offloading

FortiADC forwards data unencrypted to the servers, and the servers can maximize performance because they are processing HTTP and not HTTPS transactions.

To realize the benefits of SSL offloading and maintain security, you must deploy the FortiADC appliance in a trusted network with a direct path to the real servers so that the connection between the FortiADC and the real server does not have to be re-encrypted. For example, you connect FortiADC and the real servers through the same switch, and all are physically located on the same locked rack.

In cases where traffic is forwarded along untrusted paths toward the real servers, you can use a real server SSL profile to re-encrypt the data before forwarding it to the real servers.

Basic steps:

1. Import the X.509 v3 server certificates and their private keys that ordinarily belong to the backend servers, as well as any certificate authority (CA) or intermediate CA certificates that are used to complete the chain of trust between your clients and servers.
2. Configure a local certificate group that includes the server's local certificate and the Intermediate CA group that contains the Intermediate CAs.
3. Configure an application profile and a client SSL profile (if needed) that reference the local certificate group and specify the allowed SSL/TLS versions and list of SSL ciphers that can be used for the SSL connection between the client and the FortiADC unit. Select this profile when you configure the virtual server.

4. Configure a real server SSL profile that enables or disables SSL for the connection between the FortiADC unit and the real server. If enabled, specify the SSL/TLS versions and the list of SSL ciphers that can be used. Select this profile when you configure the real server pool.

SSL decryption by forward proxy

You can use SSL decryption by forward proxy in cases where you cannot copy the server certificate and private key to the FortiADC unit because it is either impractical or impossible (in the case of outbound traffic to unknown Internet servers).

When SSL forward proxy is enabled, FortiADC becomes a proxy to both sides of the connection. The server certificate and its private key used to negotiate the SSL connection with the client are dynamically derived from the certificate presented by the real server and optionally chained with an Intermediate CA trusted by the client.

Basic steps:

1. Import a special Intermediate CA and its private key to the local certificate store that you have provisioned for SSL forward proxy operations.
2. Configure an Intermediate CA group. (Optional)
3. Configure a certificate caching object (or use the pre-defined one).
4. Configure a client SSL profile that enables SSL proxy, references the local certificate, and specifies the allowed SSL/TLS versions and list of SSL ciphers that can be used for the SSL connection between the client and the FortiADC unit. Select this profile when you configure the virtual server.
5. Configure all settings required for backend SSL.

Layer 7 deployments

Figure 87 illustrates a Layer 7 SSL forward proxy deployment similar to the SSL offloading example—inbound traffic to your server farm. When the FortiADC virtual server receives the ClientHello message, it selects a real server and sends its own ClientHello to the server to set up its own SSL session with it (represented by the dashed line in the figure). FortiADC uses the certificate presented by the server to derive the certificate to present to the client. This derived certificate is signed by an Intermediate CA that is trusted by the client, so the client completes its handshake with the FortiADC, and FortiADC can decrypt the traffic.

Figure 87: Layer 7 SSL decryption by forward proxy

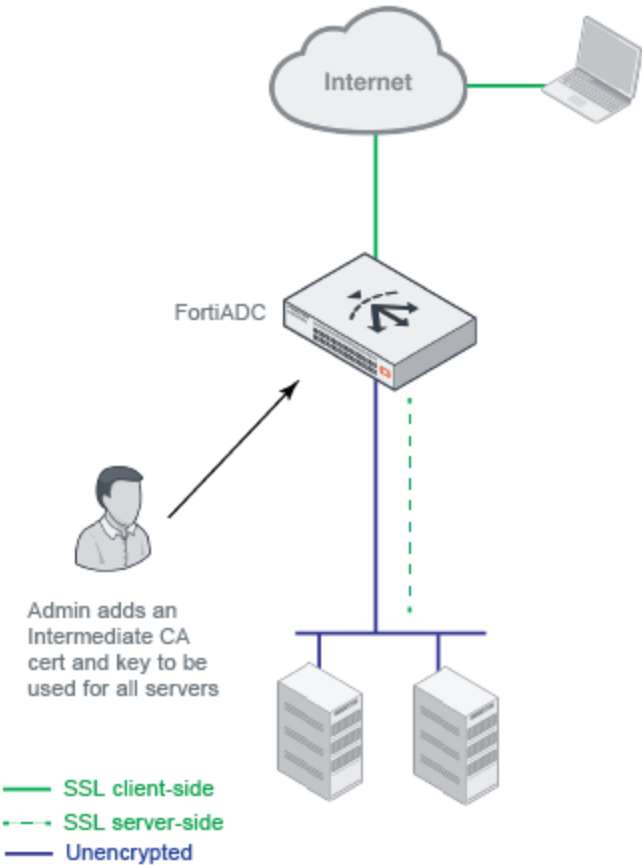


Table 161 summarizes the pros and cons of Layer 7 SSL decryption methods.

Table 161: Layer 7 SSL decryption methods

Method	Pros	Cons
SSL offloading	<p>Better performance.</p> <p>No feature limitations.</p> <p>In most cases, you do not need to maintain SSL functionality (certificates and keys, SSL ports) on the real servers.</p>	<p>You must be able to copy the local certificates and private keys from the real servers.</p>
SSL forward proxy	<p>You do not need to copy the local certificates and keys from the real servers. Instead, you add only one Intermediate CA and private key to be used for all the HTTPS servers.</p>	<p>Performance cost associate with SSL proxy operations and certificate re-signing.</p> <p>You need to maintain SSL functionality on the real servers.</p> <p>Incompatible with some features because the server must be selected before the client request is decrypted: Incompatible features include:</p> <ul style="list-style-type: none"> • Some load balancing methods (only Round Robin and Least Connection are supported) • Some persistence methods (only Source Address, Source Address Hash, Source Address-Port Hash, and SSL Session ID are supported) • Client SNI Required option • Content routing


Layer 2 deployments

You can use FortiADC in a Layer 2 sandwich topology to offload SSL decryption tasks from FortiGate.

Figure 88 shows the topology. To decrypt traffic to and from external HTTPS destinations, you must use SSL forward proxy.

When the FortiADC virtual server receives the ClientHello message, it sends its own ClientHello to the destination server in order to fetch the server certificate so that it can be manipulated. The FortiGate and second FortiADC in the network path must be configured to pass-through this HTTPS traffic. FortiADC uses the server certificate to derive a certificate to present to the client. This derived certificate is signed by an Intermediate CA that is trusted by the client, so the client completes its handshake with the first FortiADC, and FortiADC decrypts the traffic.

In a sandwich deployment like this one, you do not want to re-encrypt the traffic until it egresses the second FortiADC. You control server-side SSL with the real server SSL profile configuration, discussed next.

Figure 88: Layer 2 SSL decryption by forward proxy

Profile configurations

The application profile and client SSL profile determine the settings for the client-FortiADC connection; the real server SSL profile determines settings for the FortiADC-real server connection. This granularity gives you flexibility in how you leverage FortiADC's SSL transaction capabilities. For example, in the case of SSL

offloading, your goal is to eliminate SSL transactions on the real servers so that you can configure a server-side SSL profile that does not use SSL. Or it could be the case that the back-end real servers support only SSLv2, but you want to use the more secure TLSv1.2 for the client-FortiADC segment.


Figure 89 illustrates the basic idea of client-side and server-side profiles.

Figure 89: SSL profiles

The call-outs in Figure 90 have guidance for the two types of profiles used in a Layer 2 sandwich deployment.

In this deployment, the FortiADC 1 virtual server is of a Layer-2 HTTPS server configuration. Its client SSL profile supports SSL forward proxy, including the special local signing CA. For Layer-2 virtual servers, the "real server" target is the next hop. In this case, the real server target is the FortiGate pool. Because SSL is not enabled in the real server SSL profile, FortiADC 1 does not re-encrypt the SSL connection. (However, you can configure allowed SSL versions and ciphers in the client SSL profile, and you can also configure an SSL certificate verification policy to enforce rules and checks on the destination server certificate.) The client SSL profile settings are used when re-encrypting the server response traffic in the return segment to the client.

The FortiADC 2 virtual server is a Layer 2 HTTP virtual server configuration. It receives unencrypted traffic from FortiGate. Its server pool is the next hop gateway. On its server side, FortiADC uses the real server SSL profile settings when it encrypts the outbound SSL connection and decrypts the inbound response traffic.

Figure 90: Layer 2 sandwich profiles

For information on virtual server profile configuration objects, see [Configuring Application profiles](#).

For information on real server SSL configuration objects, see [Configuring real server SSL profiles](#).

Certificate guidelines

When a client browser requests an HTTPS connection to a web server, the server presents a server certificate to the client for verification. The client checks the content of the certificate against a local browser database of Certificate Authorities, and if it finds a match, the connection is made. If no match is found, the browser displays a warning that asks if you want to continue with the connection.

To avoid this warning, you must upload an Intermediate CA signed by one of the CA vendors that has its root certificates preinstalled in the web browsers. When the vendor issues you a local server certificate for your website, it typically includes the Intermediate CAs in your package.

For SSL offloading deployments, you create a local certificate group that references the local certificate for the server and its Intermediate CA group (a group that references all Intermediate CAs the vendor provided with your certificate package).

For SSL decryption by forward proxy deployments, you create a local certificate group that references any local certificate and an Intermediate CA group that includes the Intermediate CA and private key configuration you have provisioned for the SSL forward proxy operations.



You are not required to obtain SSL certificates from SSL vendors. You can use an enterprise certificate server (like Microsoft CertSrv) or open-source tools like OpenSSL or to generate them. Note, however, that a web browser will not trust the certificate unless it is associated with a certificate installed in the browser. If you use your own tools to generate the Intermediate CA, you must distribute that certificate to client browsers in whatever manner you typically do that—automatic update package from IT, manual distribution, and so on.

For information on importing certificates and configuring certificate configuration objects, see [Manage and validate certificates](#).

SSL/TLS versions and cipher suites

An SSL cipher is an algorithm that performs encryption and decryption. It transforms plain text into a coded set of data (cipher text) that is not reversible without a key. During the SSL handshake phase of the connection, the client sends a list of the ciphers it supports. FortiADC examines the client cipher list in the order it is specified, chooses the first cipher that matches a cipher specified in the virtual server configuration, and responds to the client. If none of the ciphers offered by the client are in the cipher suite list for the virtual server, the SSL handshake fails.

To see the list of ciphers supported by the browser you are using, go to a link maintained by the Leibniz University of Hannover Distributed Computing & Security (DCSec) Research Group:

<https://cc.dcsec.uni-hannover.de/>

FortiADC SLB profiles support a specific list of [RSA ciphers](#), [PFS ciphers](#), [ECDHE ciphers](#), [ECDSA ciphers](#), and [eNull ciphers](#).

[Table 162](#) lists supported RSA ciphers.

Table 162: Cipher suites with RSA key exchange

Abbreviation	Cipher Suite	Protocol	Kx	Au	Enc	MAC
AES256-GCM-SHA384	TLS_RSA_WITH_AES_256_GCM_SHA384	TLS 1.2	RSA	RSA	AESGCM (256)	AEAD
AES256-SHA256	TLS_RSA_WITH_AES_256_CBC_SHA256	TLS 1.2	RSA	RSA	AES(256)	SHA
AES256-SHA	TLS_RSA_WITH_AES_256_CBC_SHA	SSL 3.0 TLS 1.2, 1.1, 1.0	RSA	RSA	AES(256)	SHA
AES128-GCM-SHA256	TLS_RSA_WITH_AES_128_GCM_SHA256	TLS 1.2	RSA	RSA	AESGCM (128)	AEAD
AES128-SHA256	TLS_RSA_WITH_AES_128_CBC_SHA256	TLS 1.2	RSA	RSA	AES(128)	SHA
AES128-SHA	TLS_RSA_WITH_AES_128_CBC_SHA	SSL 3.0 TLS 1.2, 1.1, 1.0	RSA	RSA	AES(128)	SHA
RC4-SHA	SSL_RSA_WITH_RC4_128_SHA	SSL 3.0	RSA	RSA	RC4	SHA
	TLS_RSA_WITH_RC4_128_SHA	TLS 1.2, 1.1, 1.0	RSA	RSA	RC4	SHA
RC4-MD5	SSL_RSA_WITH_RC4_128_MD5	SSL 3.0	RSA	RSA	RC4	MD5
	TLS_RSA_WITH_RC4_128_MD5	TLS 1.2, 1.1, 1.0	RSA	RSA	RC4	MD5
DES-CBC3-SHA	SSL_RSA_WITH_3DES_EDE_CBC_SHA	SSL 3.0	RSA	RSA	DES-CBC3	SHA
	TLS_RSA_WITH_3DES_EDE_CBC_SHA	TLS 1.2, 1.1, 1.0	RSA	RSA	DES-CBC3	SHA
DES-CBC-SHA	SSL_RSA_WITH_DES_CBC_SHA	SSL 3.0	RSA	RSA	DES-CBC	SHA
	TLS_RSA_WITH_DES_CBC_SHA	TLS 1.2, 1.1, 1.0	RSA	RSA	DES-CBC	SHA

With RSA ciphers, the server's public RSA key is part of the server certificate and is typically very long lived. It is not uncommon for the same public key to be used for months or years. This creates a potential problem: if an SSL server's private key were to be leaked or stolen, all connections made in the past using that key would be vulnerable. If someone has recorded your SSL connections, they can use the stolen private key to decrypt them.

Table 163 lists supported Perfect Forward Secrecy (PFS) ciphers with DHE/EDH key exchange. With PFS, a fresh public key is created for every single connection. That means that an adversary would need to break the key for each connection individually to read the communication.

Table 163: Cipher suites with DHE/EDH key exchange

Abbreviation	Cipher Suite	Protocol	Kx	Au	Enc	MAC
DHE-RSA-AES256-GCM-SHA384	TLS_DHE_RSA_WITH_AES_256_GCM_SHA384	TLS 1.2	DH	RSA	AES256	SHA384
DHE-RSA-AES256-SHA256	TLS_DHE_RSA_WITH_AES_256_CBC_SHA256	TLS 1.2	DH	RSA	AES256	SHA256
DHE-RSA-AES256-SHA	TLS_DHE_RSA_WITH_AES_256_CBC_SHA	SSL 3.0 TLS 1.2, 1.1, 1.0	DH	RSA	AES256	SHA256
DHE-RSA-AES128-GCM-SHA256	TLS_DHE_RSA_WITH_AES_128_GCM_SHA256	TLS 1.2	DH	RSA	AES128	SHA256
DHE-RSA-AES128-SHA256	TLS_DHE_RSA_WITH_AES_128_CBC_SHA256	TLS 1.2	DH	RSA	AES128	SHA256
DHE-RSA-AES128-SHA	TLS_DHE_RSA_WITH_AES_128_CBC_SHA	SSL 3.0 TLS 1.2, 1.1, 1.0	DH	RSA	AES128	SHA
EDH-RSA-DES-CBC3-SHA	TLS_DHE_RSA_WITH_3DES_EDE_CBC_SHA	SSL 3.0 TLS 1.2, 1.1, 1.0	DH	RSA	3DES	SHA
EDH-RSA-DES-CBC-SHA	TLS_DHE_RSA_WITH_DES_CBC_SHA	SSL 3.0 TLS 1.2, 1.1, 1.0	DH	RSA	DES	SHA

Table 164 lists supported PFS ciphers with Elliptic curve Diffie–Hellman Ephemeral key (ECDHE) key exchange. ECDHE is significantly faster than DHE. The supported suites include both the Elliptic Curve Digital Signature Algorithm (ECDSA) and RSA key authentication (Au) algorithms.

Table 164: Cipher suites with ECDHE key exchange

Abbreviation	Cipher Suite	Protocol	Kx	Au	Enc	MAC
ECDHE-ECDSA-AES256-GCM-SHA384	TLS_ECDHE_ECDSA_WITH_AES_256_GCM_SHA384	TLS 1.2	ECDH	ECDSA	AESGCM256	AEAD
ECDHE-ECDSA-AES256-SHA384	TLS_ECDHE_ECDSA_WITH_AES_256_CBC_SHA384	TLSv1.2	ECDH	ECDSA	AES256	SHA384

Abbreviation	Cipher Suite	Protocol	Kx	Au	Enc	MAC
ECDHE-ECDSA-AES256-SHA	TLS_ECDHE_RSA_WITH_AES_256_CBC_SHA	SSL 3.0 TLS 1.2, 1.1, 1.0	ECDH	ECDSA	AES256	SHA
ECDHE-ECDSA-AES128-GCM-SHA256	TLS_ECDHE_ECDSA_WITH_AES_128_GCM_SHA256	TLSv1.2	ECDH	ECDSA	AESGCM128	AEAD
ECDHE-ECDSA-AES128-SHA256	TLS_ECDHE_ECDSA_WITH_AES_128_CBC_SHA256	TLSv1.2	ECDH	ECDSA	AES128	SHA256
ECDHE-ECDSA-AES128-SHA	TLS_ECDHE_ECDSA_WITH_AES_128_CBC_SHA	SSL 3.0 TLS 1.2, 1.1, 1.0	ECDH	ECDSA	AES128	SHA
ECDHE-ECDSA-RC4-SHA	TLS_ECDHE_ECDSA_WITH_RC4_128_SHA	SSL 3.0 TLS 1.2, 1.1, 1.0	ECDH	ECDSA	RC4	SHA
ECDHE-ECDSA-DES-CBC3-SHA	TLS_ECDHE_ECDSA_WITH_3DES_EDE_CBC_SHA	SSL 3.0 TLS 1.2, 1.1, 1.0	ECDH	ECDSA	3DES	SHA
ECDHE-RSA-AES256-GCM-SHA384	TLS_ECDHE_RSA_WITH_AES_256_GCM_SHA384	TLS 1.2	ECDH	RSA	AESGCM256	AEAD
ECDHE-RSA-AES256-SHA384	TLS_ECDHE_RSA_WITH_AES_256_CBC_SHA384	TLS 1.2	ECDH	RSA	AES256	SHA384
ECDHE-RSA-AES256-SHA	TLS_ECDHE_RSA_WITH_AES_256_CBC_SHA	TLS 1.2	ECDH	RSA	AES256	SHA
ECDHE-RSA-AES128-GCM-SHA256	TLS_ECDHE_RSA_WITH_AES_128_GCM_SHA256	TLS 1.2	ECDH	RSA	AESGCM128	AEAD
ECDHE-RSA-AES128-SHA256	TLS_ECDHE_RSA_WITH_AES_128_CBC_SHA256	TLS 1.2	ECDH	RSA	AES128	SHA256
ECDHE-RSA-AES128-SHA	TLS_ECDHE_RSA_WITH_AES_128_CBC_SHA	SSL 3.0	ECDH	RSA	AES128	SHA

Abbreviation	Cipher Suite	Protocol	Kx	Au	Enc	MAC
ECDHE-RSA-RC4-SHA	TLS_ECDHE_RSA_WITH_RC4_128_SHA	SSL 3.0	ECDH	RSA	RC4	SHA
ECDHE-RSA-DES-CBC3-SHA	TLS_ECDHE_RSA_WITH_3DES_EDE_CBC_SHA	SSL 3.0	ECDH	RSA	3DES	SHA

In addition, profiles support an eNull cipher option. This option represents all cipher suites that do not apply encryption to the application data (integrity check is still applied). The exact cipher suite used depends on the SSL/TLS version used. As an example, in SSL v3.0, eNULL includes NULL-MD5, NULL-SHA, ECDH-RSA-NUL-SHA, ECDH-ECDSA-NUL-SHA, and some other non-encryption cipher suites.

Finally, profiles support a user-specified cipher list. You can specify a colon-separated list of OpenSSL cipher suite short names. The names are validated against the form of the cipher suite short names published on the OpenSSL website:

<https://www.openssl.org/docs/manmaster/apps/ciphers.html>

Exceptions list

In some jurisdictions, SSL interception and decryption by forward proxy is disfavored for some types of websites or disallowed entirely. If necessary, you can use the L2 Exception List configuration to define destinations that should not have its sessions decrypted. You can leverage FortiGuard web filter categories, and you can configure a list of additional destinations.

You associate the L2 Exception List configuration with virtual servers that are in the path of outbound traffic. The virtual server evaluates whether an exception applies before processing the initial SSL client hello. If an exception applies, that connection is passed through, and it is not decrypted.

For information on creating the configuration, see [Configuring an L2 exception list](#).

SSL traffic mirroring

FortiADC supports mirroring packets (HTTPS/TCPS) to specified network interfaces. When the feature is enabled, SSL traffic will be mirrored to the specified ports by the virtual server after it has been decrypted.

The feature supports both IPv4 and IPv6. FortiADC can send traffic to up to four outgoing interfaces, including aggregated and VLAN interfaces. Mirrored traffic is transmitted as a single packet stream, using the original client-side source and destination IP address and port numbers. The source and destination MAC addresses are 0 (zero) in mirrored traffic. The feature requires a virtual server set to Layer 7 or Layer 2, with a profile configured for HTTPS or TCPS. It is supported on all FortiADC platforms.

Figure 91: SSL traffic mirroring

Virtual Server

Name

Status

☐ Disable
 ☒ Enable
 ☐ Maintain

Type

☒ Layer 7
 ☐ Layer 4
 ☐ Layer 2

Address Type

☒ IPv4
 ☐ IPv6

Traffic Group

default

General

Configuration

Address

Example: 192.0.2.1

Port

Default: 80 Range: 0 or 1-65535. You can enter up to 8 numbers or number ranges, e.g., 80-90 100.

Connection Limit

Default: 10000 Range: 0-100000000 concurrent connections

Interface

port1

Resources

Profile

LB_PROF_TCPS

Persistence

Click to select.

Method

LB_METHOD_DYNAMIC_LOAD

Real Server Pool

test1

SSL Traffic Mirror

SSL Traffic Mirror

☒ Enable

Mirror To

Selected Items

port2

Double-click to deselect. Drag to reorder.

Available Items

port1

port3

port4

port5

Double-click to select.

Note that this feature is available via the CLI only, and has not yet been implemented on the GUI.

To enable this feature in a policy, execute the following command:

```
config load-balance virtual-server
edit vs-name
set ssl-mirror enable
set ssl-mirror-intf port1 port2
next
end
```

FortiADC Deployment Guide
Fortinet Technologies, Inc.

430

Chapter 17: Advanced Networking

This chapter includes the following topics:

- NAT
- "Configure source NAT" on page 431
- "Configure 1-to-1 NAT" on page 434
- Configuring the QoS filter
- ISP routes
- BGP
- OSPF
- IPv4 access list
- IPv6 access list
- IPv4 prefix list
- IPv6 prefix list
- TCP multiplexing
- Reverse path route caching

NAT

A number of network address translation (NAT) methods map packet IP address information for the packets that are received at the ingress network interface into the IP address space you configure. Packets with the new IP address are forwarded through the egress interface.

You can configure NAT per virtual server within the virtual server configuration.

This section describes the system-wide, policy-based NAT feature. The system-wide feature supports:

- SNAT—Translates the packet header source IP address to the configured address. See [Configure source NAT](#).
- 1-to-1 NAT—Maps the public IP address for an interface to an IP address on a private network. See [Configure 1-to-1 NAT](#).
- Port forwarding—Maps an external published protocol port to the actual port. Configuration for port forwarding is included in the configuration for 1-to-1 NAT.

Configure source NAT

You use source NAT (SNAT) when clients have IP addresses from private networks. This ensures you do not have multiple sessions from different clients with source IP 192.168.1.1, for example. Or, you can map all client traffic to a single source IP address because a source address from a private network is not meaningful to the FortiADC system or backend servers.

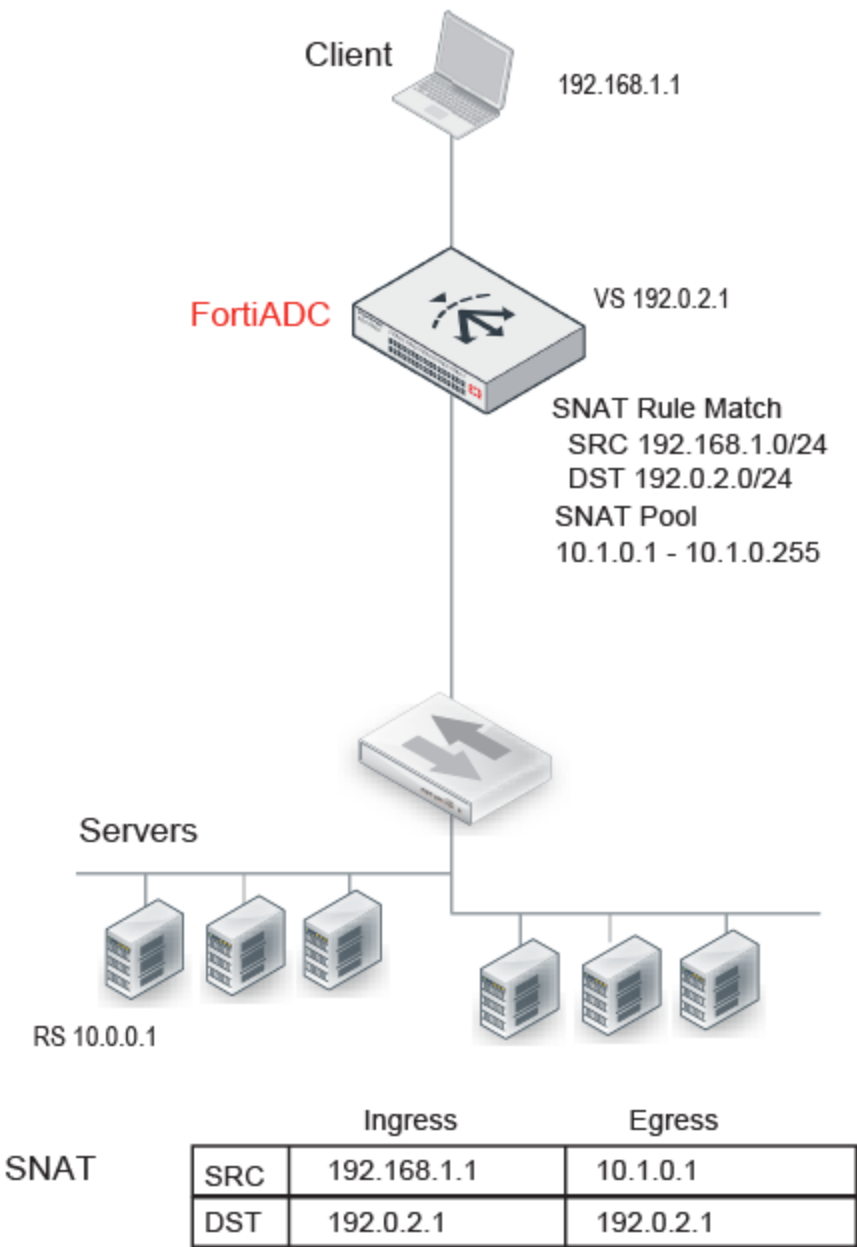
[Figure 92](#) illustrates SNAT. The SNAT rule matches the source and destination IP addresses in incoming traffic to the ranges specified in the policy. If the client request matches, the system translates the source IP address to an address from the SNAT pool. In this example, a client with private address 192.168.1.1 requests a resource from the virtual server address at 192.0.2.1 (not the real server address 10.0.0.1; the real server address is not

published). The two rule conditions match, so the system translates the source IP to the next address in the SNAT pool—10.1.0.1. SNAT rules do not affect destination addresses, so the destination address in the request packet is preserved.

The system maintains this NAT table and performs the inverse translation when it receives the server-to-client traffic. Be sure to configure the backend servers to use the FortiADC address as the default gateway so that server responses are also rewritten by the NAT module.

Note: This SNAT feature is not supported for traffic to virtual servers. Use the virtual server SNAT feature instead.

Figure 92: SNAT




Before you begin:

- You must know the IP addresses your organization has provisioned for your NAT design.
- You must have Read-Write permission for System settings.

To configure source NAT:

1. Go to Networking > NAT.
The configuration page displays the Source tab.
2. Click **Add** to display the configuration editor.
3. Complete the configuration as described in [Table 165](#).
4. Save the configuration.
5. Reorder rules, as necessary.

Table 165: Source NAT configuration

Settings	Guidelines
Name	Configuration name. Valid characters are A-Z, a-z, 0-9, _, and -. No spaces. After you initially save the configuration, you cannot edit the name.
Source	Address/mask notation to match the source IP address in the packet header. For example, 192.0.2.0/24.
Destination	Address/mask notation to match the destination IP address in the packet header. For example, 10.0.2.0/24.
Egress Interface	Interface that forwards traffic.
Translation Type	<ul style="list-style-type: none"> • IP Address—Select to translate the source IP to a single specified address. • Pool—Select to translate the source IP to the next address in a pool.
Translation to IP Address	<p>Note: This option applies only when the Translation Type is set to IP address.</p> <p>Specify an IPv4 address. The source IP address in the packet header will be translated to this address.</p>
Pool Address Range	<p>Note: This option applies only when Translation Type is set to Pool.</p> <p>Specify the first IP address in the SNAT pool.</p>
To	Specify the last IP address in the SNAT pool.
Traffic Group	Select a traffic group. Otherwise, the system will use the default traffic group.
Reordering	
	After you have saved a rule, reorder rules as necessary. The rules table is consulted from top to bottom. The first rule that matches is applied and subsequent rules are not evaluated.

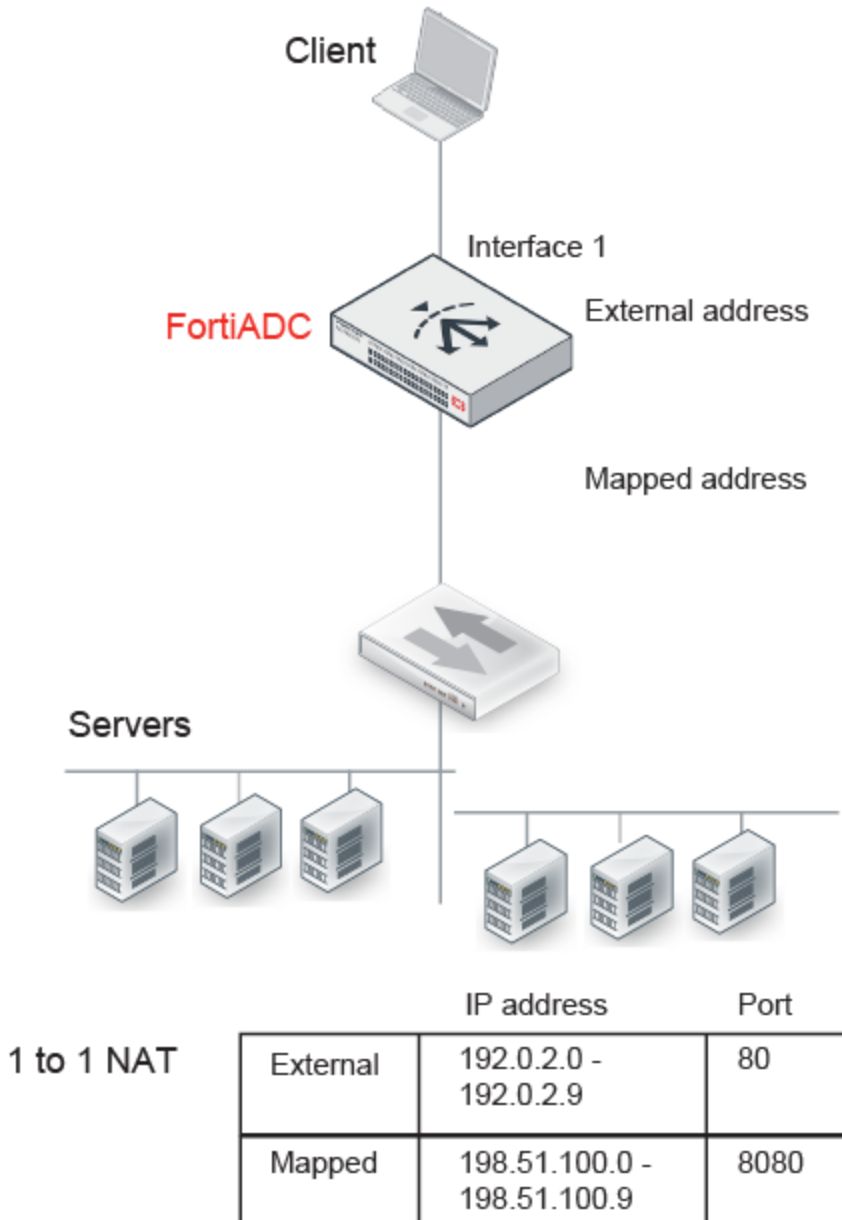
Configure 1-to-1 NAT

You can use 1-to-1 NAT when you want to publish public or “external” IP addresses for FortiADC resources but want the communication among servers on the internal network to be on a private or “internal” IP address range.

[Figure 93](#) illustrates 1-to-1 NAT. The NAT configuration assigns both external and internal (or “mapped”) IP addresses to Interface 1. Traffic from the external side of the connection (such as client traffic) uses the external IP address and port. Traffic on the internal side (such as the virtual server communication with real servers) uses the mapped IP address and port.

1-to-1 NAT is supported for traffic to virtual servers. The address translation occurs before the ADC has processed its rules, so FortiADC server load balancing policies that match source address (such as content routing and content rewriting rules) should be based on the mapped address space.

The system maintains this NAT table and performs the inverse mapping when it sends traffic from the internal side to the external side.

Figure 93: One-to-One NAT

Before you begin:


- You must know the IP addresses your organization has provisioned for your NAT design.
- You must have Read-Write permission for System settings.

To configure one-to-one NAT:

1. Go to Networking > NAT.
2. Click the **1-to-1 NAT** tab.
3. Click **Add** to display the configuration editor.
4. Complete the configuration as described in [Table 166](#).

5. Save the configuration.
6. Reorder rules, as necessary.

Table 166: 1-to-1 NAT configuration

Settings	Guidelines
Name	Configuration name. Valid characters are A-Z, a-z, 0-9, _, and -. No spaces. After you initially save the configuration, you cannot edit the name.
External Interface	Interface that receives traffic.
External Address Range	Specify the first address in the range. The last address is calculated after you enter the mapped IP range.
Mapped Address Range	Specify the first and last addresses in the range.
Port Forwarding	
Port Forwarding	Select to enable.
Protocol	<ul style="list-style-type: none"> • TCP • UDP
External Port Range	Specify the first port number in the range. The last port number is calculated after you enter the mapped port range.
Mapped Port Range	Specify the first and last port numbers in the range.
Traffic Group	Select a traffic group. Otherwise, the system will use the default.
Reordering	
	After you have saved a rule, reorder rules as necessary. The rules table is consulted from top to bottom. The first rule that matches is applied and subsequent rules are not evaluated.

QoS

You can use quality-of-service (QoS) policies to provision bandwidth for any traffic that matches the rule. You might consider QoS policies for latency- or bandwidth-sensitive services, such as VoIP and ICMP.

The FortiADC system does not provision bandwidth based on the TOS bits (also called differentiated services) in the IP header to control packet queueing. Instead, the system provisions bandwidth based on a source/destination/service matching tuple that you specify.

Note: The QoS policy feature is not supported for traffic to virtual servers.

Basic steps

1. Configure a [QoS queue](#).
2. Configure a [QoS filter](#) or [QoS IPv6 filter](#).

Configuring a QoS queue

You must configure a queue before you configure a filter.

Before you begin:

- You must have Read-Write permission for System settings.

To configure a QoS queue:

1. Go to Networking > QoS.
2. Click the **QoS Queue** tab.
3. Click **Add** to display the configuration editor.
4. Complete the configuration as described in [Table 167](#)
5. Save the configuration.

Table 167: QoS queue configuration

Settings	Guidelines
Name	Configuration name. Valid characters are A-Z, a-z, 0-9, _, and -. No spaces. After you initially save the configuration, you cannot edit the name.
Bandwidth	Maximum bandwidth rate. Specify a number and a unit abbreviation. For example, specify 100K for 100 Kbps, 10M for 10 Mbps, and 1G for 1Gbps.

Configuring the QoS filter

A QoS filter is the policy that assigns traffic to the QoS queue.

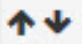
Before you begin:

- You must have a good understanding and knowledge of traffic in your network that requires QoS provisioning.
- You must have created the address configuration objects and service configuration objects that define the matching tuple for QoS rules. Use the Shared Resources menu firewall address and service object configuration editor.
- You must have created a QoS queue configuration object.
- You must have Read-Write permission for System settings.

To configure QoS filter:

1. Go to Networking > QoS.
2. Click the **QoS Filter** tab.
3. Click **Add** to display the configuration editor.
4. Complete the configuration as described in [Table 168](#).
5. Save the configuration.
6. Reorder rules, as necessary.

Table 168: QoS filter configuration

Settings	Guidelines
Name	Configuration name. Valid characters are A-Z, a-z, 0-9, _, and -. No spaces. After you initially save the configuration, you cannot edit the name.
Status	Enable/disable the filter.
Queue	Select the queue that will be used for packets that match the filter criteria.
Service	Select a service object to use to form the matching tuple.
Source	Select a source address object to use to form the matching tuple.
Destination	Select a destination address object to use to form the matching tuple.
Ingress Interface	Select the interface that receives traffic.
Egress Interface	Select the interface that forwards traffic.
	After you have saved a rule, reorder rules as necessary. The rules table is consulted from top to bottom. The first rule that matches is applied and subsequent rules are not evaluated.

Configuring the QoS IPv6 filter

A QoS filter is the policy that assigns traffic to the QoS queue.

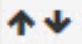
Before you begin:

- You must have a good understanding and knowledge of traffic in your network that requires QoS provisioning.
- You must have created the address configuration objects and service configuration objects that define the matching tuple for QoS rules. Use the Shared Resources menu firewall address and service object configuration editor.
- You must have created a QoS queue configuration object.
- You must have Read-Write permission for System settings.

To configure QoS filter:

1. Go to Networking > QoS.
2. Click the **QoS IPv6 Filter** tab.
3. Click **Add** to display the configuration editor.
4. Complete the configuration as described in [Table 169](#).
5. Save the configuration.
6. Reorder rules, as necessary.

Table 169: QoS IPv6 filter configuration

Settings	Guidelines
Name	Configuration name. Valid characters are A-Z, a-z, 0-9, _, and -. No spaces. After you initially save the configuration, you cannot edit the name.
Status	Enable/disable the filter.
Queue	Select the queue that will be used for packets that match the filter criteria.
Service	Select a service object to use to form the matching tuple.
Source	Select a source address object to use to form the matching tuple.
Destination	Select a destination address object to use to form the matching tuple.
Ingress Interface	Select the interface that receives traffic.
Egress Interface	Select the interface that forwards traffic.
	After you have saved a rule, reorder rules as necessary. The rules table is consulted from top to bottom. The first rule that matches is applied and subsequent rules are not evaluated.

ISP routes

ISP routes can be used for outbound traffic and link load balancing traffic.

Routes for outbound traffic are chosen according to the following priorities:

1. Link local routes—Self-traffic uses link local routes.
2. LLB Link Policy route—Configured policy routes have priority over default routes.
3. Policy route—Configured policy routes have priority over default routes.
4. Static route / ISP route / OSPF route—Priority is based on the distance metric. By default, distance for static routes is 10, for ISP routes is 20, and for OSPF routes is 110. The distance metric is configurable for static routes and OSPF routes, but not ISP routes.
5. Default LLB Link Policy route—Default routes have lower priority than configured routes.
6. Default static route / OSPF route—Default routes have lower priority than configured routes.

Before you begin:

- You must have read-write permission for system settings.

Note: Adding a new ISP route does not affect existing sessions. Deleting or editing an ISP route causes the related sessions to be re-created.

To configure ISP Routes:

1. Go to Networking > Routing.
2. Click the **ISP** tab.
3. Click **Add** to display the configuration editor.

4. Complete the configuration as described in [Table 170](#).
5. Save the configuration.

Table 170: ISP Route configuration

Settings	Guidelines
Destination	Select an ISP address book configuration object. Note: Two ISP routes cannot reference the same ISP address book. The ISP routing feature does not support multipath routing.
Gateway	IP address of the gateway router that can route packets to the destination IP address that you have specified.

BGP

BGP stands for Border Gateway Protocol, which was first used in 1989. The current version, BGP-4, was released in 1995 and is defined in RFC 1771. That RFC has since been replaced by the more recent RFC 4271. The main benefits of BGP-4 are classless inter-domain routing and aggregate routes. Often classified as a path-vector protocol and sometimes as a distance-vector routing protocol, BGP exchanges routing and reachability information among autonomous systems over the Internet.

BGP makes routing decisions based on path, network policies and rulesets instead of the hop-count metric as RIP does, or cost-factor metrics as OSPF does.

BGP-4+ supports IPv6. It was introduced in RFC 2858 and RFC 2545.

BGP is the routing protocol used on the Internet. It was designed to replace the old Exterior Gateway Protocol (EGP) which had been around since 1982, and was very limited. In doing so, BGP enabled more networks to take part in the Internet backbone to effectively decentralize it and make the Internet more robust, and less dependent on a single ISP or backbone network.

How BGP works

A BGP router receives information from its peer routers that have been defined as neighbors. BGP routers listen for updates from these configured neighboring routers on TCP port 179.

A BGP router is a finite state machine with six various states for each connection. As two BGP routers discover each other, and establish a connection they go from the idle state, through the various states until they reach the established state. An error can cause the connection to be dropped and the state of the router to be reset to either active or idle. These errors can be caused by: TCP port 179 not being open, a random TCP port above port 1023 not being open, the peer address being incorrect, or the AS number being incorrect.

When BGP routers start a connection, they negotiate which (if any) optional features will be used such as multiprotocol extensions that can include IPv6 and VPNs.

IBGP vs. EBGP

When you read about BGP, often you see EBGP or IBGP mentioned. These are both BGP routing, but BGP used in different roles. Exterior BGP (EBGP) involves packets crossing multiple autonomous systems (ASes) where

interior BGP (IBGP) involves packets that stay within a single AS. For example the AS_PATH attribute is only useful for EBGp where routes pass through multiple ASes.

These two modes are important because some features of BGP are only used for one of EBGp or IBGP. For example confederations are used in EBGp, and route reflectors are only used in IBGP. Also routes learned from IBGP have priority over EBGp learned routes.

For more information on BGP routing, see "Chapter 3 - Advanced Routing" of the *FortiOS Handbook for FortiOS 5.4.1*.

Before you begin, you must:

- Know how BGP has been implemented in your network, i.e., the configuration details of the implementation.
- Have Read-Write permission for System settings.
- Have configured all the needed access (IPv6) lists and prefix (IPv6) lists. See [Access list vs. prefix list](#).

To configure BGP:

1. Click Networking > Routing.
2. Click the **BGP** tab.
3. Make the desired entries and/or selections as described in the table below.
4. Click **Save** when done.

Table 171: BGP configuration

Settings	Guidelines
AS	<p>Enter the AS (Autonomous System) number of the BGP router. Valid values are from 0 to 4294967295.</p> <p>Note: Per RFC 6996, the first and last ASNs of the original 16-bit integers, namely 0 and 65535, and the last ASN of the 32-bit numbers, namely 4,294,967,295, are reserved and should not be used by operators; ASNs 64,512 to 65,534 of the original 16-bit AS range, and 4,200,000,000 to 4,294,967,294 of the 32-bit range are reserved for private use, which means that they can be used internally but should not be announced to the global Internet.</p>
Router ID	Enter the 32-bit number that sets the router-ID of the BGP process. The router ID uses dotted decimal notation. The router-ID must be the IP address of the router, and it must be unique within the entire BGP domain to the BGP speaker.
Redistribute OSPF	Enable/Disable (default) the redistribution of OSPF routes to the BGP process.
Redistribute Connected	Enable/Disable (default) the redistribution of connected routes to the BGP process.
Redistribute Static	Enable/Disable (default) the redistribution of static routes to the BGP process.
Redistribute IPv6 Connected	Enable/Disable (default) the redistribution of connected IPv6 routes to the BGP process.

Settings	Guidelines
Redistribute IPv6 Static	Enable/Disable (default) the redistribution of static IPv6 routes to the BGP process.
Always Compare MED	Enable/Disable (default) the comparison of Multi-Exit Discriminator (MED) for paths from neighbors in different ASs (Autonomous Systems).
Deterministic MED	Enable/Disable (default) the deterministic comparison of Multi-Exit Discriminator (MED) values among all paths received from the same AS (Autonomous System).
Bestpath Compare Router ID	Enable/Disable (default) the BGP routing process to compare identical routes received from different external peers during the best-path selection process and to select the route with the lowest router ID as the best path.
Network	
Type	Select either of the following (IP address) types: <ul style="list-style-type: none"> • IPv4 • IPv6
IPv4 Prefix	If IPv4 is selected (above), specify the IPv4 prefix in the format of 0.0.0.0/0.
IPv6 Prefix	If IPv6 is selected (above), specify the IPv6 prefix in the format of ::/0.
Save	Be sure to click Save after you are done with configuring the network.
Neighbor	
Remote AS	Specify the remote AS (Autonomous System) number of the BGP neighbor you are creating. Valid values are from 1 to 4294967295.
Type	Select either of the following: <ul style="list-style-type: none"> • IPv4 • IPv6
IP/IPv6	Specify the IPv4 address or IPv6 address for the BGP neighbor.
Interface	Click to select the interface for the BGP neighbor.
Port	Specify the port of the BGP neighbor.
Keep Alive	Specify the frequency (in seconds) at which the BGP neighbor sends out <i>keepalive</i> message to its peer. Valid values are from 0 to 65535, with 60 seconds being the default.

Settings	Guidelines
Hold Time	<p>Specify the "wait time" or pause (in seconds) the BGP neighbor declares a peer dead after failing to receive a <i>keepalive</i> message from it.</p> <p>Valid values are from 0 to 65535, with 180 (seconds) being the default.</p> <p>When the minimum acceptable hold time is configured on a BGP router, a remote BGP peer session can be established only when the latter is advertising a hold time equal to, or greater than, the minimum acceptable hold time configured on the former. If the minimum acceptable hold time is greater than the configured hold time, then the next time the remote BGP peer tries to establish a session with the local BGP router, it will fail and the local BGP router will notify the remote BGP peer saying "unacceptable hold time".</p>
Distribute List In/Distribute IPv6 List In	<p>Click to select an Access list or Access IPv6 List.</p> <p>The BGP router will apply the selected prefix list to inbound advertisements to the BGP neighbor when distributing BGP neighbor information.</p> <p>Note: It is highly recommended that you have the IPv4 Prefix List or the IPv6 Prefix List configured before configuring BGP Routing.</p>
Distribute List Out/Distribute IPv6 List Out	<p>Click to select an Access List or Access IPv6 List.</p> <p>The BGP router will apply the selected prefix list to outbound advertisements to the neighbor when distributing BGP neighbor information.</p> <p>Note: It is highly recommended that you have the IPv4 Prefix List or the IPv6 Prefix List configured before configuring BGP Routing.</p>
Prefix List In/Prefix IPv6 List In	<p>Click to select an Prefix List (for IPv4) or Prefix IPv6 List (for IPv6).</p> <p>The BGP router will apply the selected prefix list to inbound advertisements to the neighbor when distributing BGP neighbor information.</p> <p>Note: It is highly recommended that you have the Prefix List or the Prefix IPv6 List configured before configuring BGP Routing.</p>
Prefix List Out/Prefix IPv6 List Out	<p>Click to select an IPv4 Prefix List or IPv6 Prefix List.</p> <p>The BGP router will apply the selected prefix list to outbound advertisements to the neighbor when distributing BGP neighbor information.</p> <p>Note: It is highly recommended that you have the IPv4 Prefix List or the IPv6 Prefix List configured before configuring BGP Routing.</p>
Weight	<p>Assign a weight to a neighbor connection. Valid values are from 0 to 65535.</p> <p>By default, routes learned through another BGP peer carries a weight value of 0, whereas routes sourced by the local router carry a default weight value of 32768.</p> <p>Initially, all routes learned from a neighbor will have an assigned weight. The route with the greatest weight is chosen as the preferred route when multiple routes are available to a network.</p>

Settings	Guidelines
Save	Be sure to click Save after you are done with configuring the Neighbor.
HA Router ID List	
Router ID	<p>Use the HA Router list configuration in an HA active-active deployment. On each HA cluster node, add an HA Router configuration that includes an entry for each cluster node. When the appliance is in standalone mode, it uses the primary BGP Router ID; when it is in HA mode, it uses the HA Router list ID.</p> <p>Specify a 32-bit number that sets the router-ID of the BGP process. The router ID uses dotted decimal notation. The router-ID must be an IP address of the router, and it must be unique within the entire BGP domain to the BGP speaker.</p>
Node	Specify the HA Node ID (0-7).
Save	Be sure to click Save after you are done with configuring the HA Router ID List.

Note: The access list and prefix list features are mutually exclusive. Therefore, do NOT try to apply both to any neighbor in any direction (inbound or outbound) when configuring BGP routing.

Access list vs. prefix list

Access lists and prefix lists are different mechanisms that you can use to control traffic into and out of a network.

Access lists

Access lists allow you to filter packets so that you can permit or deny them from crossing specified network interfaces. You can control whether packets are forwarded or blocked at the routers' interfaces based on the criteria set in the access lists.

Access lists fall into two categories: standard and extended. A standard access list (1-99) only checks the source addresses of all IP packets, whereas an extended access list (100-199) checks both source and destination addresses, specific UDP/TCP/IP protocols, and destination ports.

Table 172 below provides a comparison between standard access lists and extended access lists in terms of range.

Table 172: Range comparison between standard access list and extended access list

Access List Type	Range
Standard	1-99, 1300-1999
Extended	100-199, 2000-2699

Note: For this release, FortiADC only supports user-defined access lists. It does NOT support either standard or extended access lists. Access lists are NOT required for BGP routing configuration. However, if you want to

include access lists in BGP routing configuration, we highly recommend that you have them configured ahead of time.

Prefix list

Prefix lists are used to configure filter IP routes. They are configured with the permit or deny keywords to either allow or block the prefix based on the matching conditions. A prefix list is made up of an IP address and a bit mask. The IP address can be a classful network, a subnet, or a single host route, whereas the bit mask can be a numeric value ranging from 1 to 32. An implicit deny is applied to the route that matches any entry in the prefix list.

A prefix list contains one or multiple sequential entries which are evaluated sequentially, starting with the entry with the lowest sequence number. Evaluation of a prefix against a prefix list comes to an end when a match is found and the permit or deny statement is applied to that network.

Although extended access lists, and, to some extent, standard access lists, can be utilized to match prefix announcements, prefix lists are considered more graceful.

Note: Prefix lists are NOT required for BGP routing configuration. However, if you want to include prefix lists in BGP routing configuration, we highly recommend that you have them configured ahead of time.

Configuring an Access List

FortiADC D-Series units support IPv4 access lists over BGP routing. If you are configuring BGP routing using IPv4, you must configure access lists using the IPv4 protocol.

To configure an Access List:

1. Click Networking > Routing.
2. Click the Access List tab.
3. Click **Add**.
4. Enter a unique name for the new access list. **Note:** The name can be up to 35 alphanumeric characters long, including . (period), : (colon), _ (underscore), and - (hyphen). No space is allowed.
5. Enter a brief description of the access list. **Note:** The description can be up to 1023 alphanumeric characters long, with no restriction on use of special characters. Space between characters is allowed.
6. Click **Save**. The newly created access list entry appears in the access list table.
7. Click the Edit button to open the Access List dialog.
8. In the Rule pane, click **Add**. The Access List > Edit Rule tab opens.
9. For Action, select the **Permit** or **Deny** radio button.
10. For IPv4 Prefix, enter the IPv4 address/subnet mask in the format of 0.0.0.0/0.
11. Click **Save** when done.
12. Repeat Steps 8 through 11 above to add as many rules to the access list as needed.
13. Click **X** to close the Access List dialog when done.

Configuring an Access IPv6 List

FortiADC D-Series units support IPv6 access lists over BGP routing. If you are configuring BGP routing using IPv6, you must configure access lists using the IPv6 protocol.

To configure an Access IPv6 List:

1. Go to Network > Routing.
2. Click the Access IPv6 List tab.
3. Click **Add**.
4. Enter a unique name for the new access list. **Note:** The name can be up to 35 alphanumeric characters long, including . (period), : (colon), _ (underscore), and - (hyphen). No space is allowed.
5. Enter a brief description of the access list. **Note:** The description can be up to 1023 alphanumeric characters long, with no restriction on use of special characters. Space between characters is allowed.
6. Click **Save**. The newly created access list entry appears in the access list table.
7. Click the Edit button to open the Access IPv6 List dialog.
8. In the Rule pane, click **Add**. The Access IPv6 List > Edit Rule tab pens.
9. For Action, select the **Permit** or **Deny** radio button.
10. For IPv6 Prefix, enter the IPv6 address/subnet mask in the format of ::/0.
11. Click **Save** when done.
12. Repeat Steps 8 through 11 above to add as many rules to the access list as needed.
13. Click **X** to close the Access IPv6 List dialog when done.

Configuring a Prefix List

FortiADC D-Series units support IPv4 prefix lists over BGP routing. If you are configuring BGP routing using IPv4, you must configure access lists using the IPv4 protocol.

To configure a Prefix List:

1. Go to Network > Routing.
2. Click the Prefix List tab.
3. Click **Add**.
4. Enter a unique name for the new access list. **Note:** The name can be up to 35 alphanumeric characters long, including . (period), : (colon), _ (underscore), and - (hyphen). No space is allowed.
5. Enter a brief description of the access list. **Note:** The description can be up to 1023 alphanumeric characters long, with no restriction on use of special characters. Space between characters is allowed.
6. Click **Save**. The newly created access list entry appears in the access list table.
7. Click the Edit button to open the Prefix List dialog.
8. In the Rule pane, click **Add**. The Prefix List > Edit Rule tab pens.
9. For Action, select the **Permit** or **Deny** radio button.
10. For IPv4 Prefix, enter the IPv4 address/subnet mask in the format of 0.0.0.0/0.
11. For **GE**, set the **GE** (greater than and equal to) values.
12. For **LE**, set the **LE** (less than and equal to) values.
13. Click **Save** when done.
14. Repeat Steps 8 through 13 above to add as many rules to the access list as needed.
15. Click **X** to close the Prefix List dialog when done.

Configuring an IPv6 prefix list

FortiADC D-Series units support IPv6 prefix lists over BGP routing. If you are configuring BGP routing using IPv6, you must configure access lists using the IPv6 protocol.

To configure a Prefix IPv6 List:

1. Go to Network > Routing.
2. Click the Prefix IPv6 List tab.
3. Click **Add**.
4. Enter a unique name for the new access list. **Note:** The name can be up to 35 alphanumeric characters long, including . (period), : (colon), _ (underscore), and - (hyphen). No space is allowed.
5. Enter a brief description of the access list. **Note:** The description can be up to 1023 alphanumeric characters long, with no restriction on use of special characters. Space between characters is allowed.
6. Click **Save**. The newly created access list entry appears in the access list table.
7. Click the Edit button to open the Prefix IPv6 List dialog.
8. In the Rule pane, click **Add**. The Prefix IPv6 List > Edit Rule tab pens.
9. For Action, select the **Permit** or **Deny** radio button.
10. For IPv6 Prefix, enter the IPv6 address/subnet mask in the format of ::/0.
11. For **GE**, set the **GE** (greater than and equal to) values.
12. For **LE**, set the **LE** (less than and equal to) values
13. Click **Save** when done.
14. Repeat Steps 8 through 13 above to add as many rules to the access list as needed.
15. Click **X** to close the Prefix IPv6 List dialog when done.

OSPF

OSPF (Open Shortest Path First) is described in RFC2328, OSPF Version 2. It is a link-state interior routing protocol. Compared with RIP, OSPF can provide scalable network support and faster convergence times. OSPF is widely used in large networks such as ISP backbone and enterprise networks. FortiADC supports OSPF version 2.

Before you begin:

- You must know how OSPF has been implemented in your network, and you must know the configuration details of the implementation.
- You must have Read-Write permission for System settings.

To configure OSPF:

1. Go to Networking > Routing.
2. Click the **OSPF** tab.
3. Click **Add** to display the configuration editor.
4. Complete the configuration as described in [Table 173](#).
5. Save the configuration.

Table 173: OSPF configuration

Settings	Guidelines
Router	A 32-bit number that sets the router-ID of the OSPF process. The router ID uses dotted decimal notation. The router-ID must be an IP address of the router, and it must be unique within the entire OSPF domain to the OSPF speaker.
Default Metric	The default is 10.
Distance	The default is 110.
Default Information Originate	<ul style="list-style-type: none"> • Disable—Default. • Enable—Originate an AS-External (type-5) LSA describing a default route into all external routing capable areas of the specified metric and metric type. • Always—The default is always advertised even when there is no default route present in the routing table.
Default Information Metric	The default is -1, which equals to the Default Metric.
Default Information Metric Type	Select either of the following: <ul style="list-style-type: none"> • 1—If selected, the metric equals to the Default Information Metric, plus the Default Metric. • 2—(Default) If selected, the metric equals to the Default Information Metric.
Redistribute Connected	Enable/disable to redistribute connected routes to OSPF, with the metric type and metric set if specified. Redistributed routes are distributed to OSPF as Type-5 External LSAs into links to areas.
Redistribute Connected Metric	The default is -1, which equals to the Default Metric.
Redistribute Connected Metric Type	Select either of the following: <ul style="list-style-type: none"> • 1—If selected, the metric equals to the Redistribute Connected Metric, plus the Default Metric. • 2—(Default) If selected, the metric equals to the Redistribute Connected Metric.
Redistribute Static	Enable/disable to redistribute static routes into OSPF, with the metric type and metric set if specified. Redistributed routes are distributed into OSPF as Type-5 External LSAs into links to areas.
Redistribute Static Metric	The default is -1, which equals to the Default Metric.
Redistribute Static Metric Type	<ul style="list-style-type: none"> • 1—If selected, the metric equals to the Redistribute Static Metric, plus the Default Metric. • 2—(Default) If selected, the metric equals to the Redistribute Static Metric.
Area Authentication	


Settings	Guidelines
Area	A 32-bit number that identifies the OSPF area. An OSPF area is a smaller part of the larger OSPF network. Areas are used to limit the link-state updates that are sent out. The flooding used for these updates would overwhelm a large network, so it is divided into these smaller areas for manageability.
Authentication	<p>Specify an authentication type:</p> <ul style="list-style-type: none"> • None—Also called null authentication. No authentication is used. In this case the 16-byte Authentication field is not checked, and can be any value. However checksumming is still used to locate errors. • Text—A simple password is used. The password is a plain text string of characters. The same password is used for all transactions on a network. The main use of this type of authentication is to prevent routers from accidentally joining the network. Simple password authentication is vulnerable to many forms of attack, and is not recommended as a secure form of authentication. • MD5—Use OSPF cryptographic authentication. A shared secret key is used to authenticate all router traffic on a network. The key is never sent over the network in the clear—a packet is sent and a condensed and encrypted form of the packet is appended to the end of the packet. A non-repeating sequence number is included in the OSPF packet to protect against replay attacks that could try to use already sent packets to disrupt the network. When a packet is accepted as authentic, the authentication sequence number is set to the packet sequence number. If a replay attack is attempted, the packet sent will be out of sequence and ignored.
Network	
Prefix	Address/mask notation to specify the subnet.
Area	Select an area configuration.
Interface	
Name	Configuration name. Valid characters are A-Z, a-z, 0-9, _, and -. No spaces. After you initially save the configuration, you cannot edit the name.
Interface	Select the interface to enable OSPF for it.
Ignore MTU	Enable/disable to ignore the interface MTU. Disabled by default.
Network Type	<ul style="list-style-type: none"> • Broadcast • Point to Point • Point to Multipoint
Retransmit Interval	Interval for retransmitting Database Description and Link State Request packets. The default is 5 seconds.
Transmit Delay	Increment LSA age by this value when transmitting. The default is 1 second.

Settings	Guidelines
Cost	Set link cost for the specified interface. The cost value is set to router-LSA's metric field and used for SPF calculation. The default is 0.
Priority	The router with the highest priority will be more eligible to become Designated Router. Setting the value to 0 makes the router ineligible to become Designated Router. The default is 1.
Dead Interval	Number of seconds for RouterDeadInterval timer value used for Wait Timer and Inactivity Timer. This value must be the same for all routers attached to a common network. The default is 40 seconds.
Hello Interval	Number of seconds between hello packets sent on the configured interface. This value must be the same for all routers attached to a common network. The default is 10 seconds.
Authentication	<p>Specify an authentication type. All OSPF interfaces that want to learn routes from each other must be configured with the same authentication type and password or MD5 key (one match is enough). Options are:</p> <ul style="list-style-type: none"> • None—Also called null authentication. No authentication is used. In this case the 16-byte Authentication field is not checked, and can be any value. However checksumming is still used to locate errors. • Text—A simple password is used. The password is a plain text string of characters. The same password is used for all transactions on a network. The main use of this type of authentication is to prevent routers from accidentally joining the network. Simple password authentication is vulnerable to many forms of attack, and is not recommended as a secure form of authentication. • MD5—Use OSPF cryptographic authentication. A shared secret key is used to authenticate all router traffic on a network. The key is never sent over the network in the clear—a packet is sent and a condensed and encrypted form of the packet is appended to the end of the packet. A non-repeating sequence number is included in the OSPF packet to protect against replay attacks that could try to use already sent packets to disrupt the network. When a packet is accepted as authentic, the authentication sequence number is set to the packet sequence number. If a replay attack is attempted, the packet sent will be out of sequence and ignored.
Text	If using text authentication, specify a password string. Passwords are limited to 8 characters.
MD5	If using MD5 authentication, select an MD5 configuration name.
HA Router	

Settings	Guidelines
Router	<p>You use the HA Router list configuration in an HA active-active deployment. On each HA cluster node, add an HA Router configuration that includes an entry for each cluster node. When the appliance is in standalone mode, it uses the primary OSPF Router ID; when it is in HA mode, it uses the HA Router list ID.</p> <p>Specify a 32-bit number that sets the router-ID of the OSPF process. The router ID uses dotted decimal notation. The router-ID must be an IP address of the router, and it must be unique within the entire OSPF domain to the OSPF speaker.</p>
Node	HA Node ID (0-7).
MD5 Key List	
Name	<p>Configuration name. You select this name in the OSPF Interface configuration.</p> <p>Valid characters are A-Z, a-z, 0-9, _, and -. No spaces. After you initially save the configuration, you cannot edit the name.</p>
Member	
Key ID	A number that identifies a member. Valid values are from 1 to 255. Each member key ID must be unique to its member list.
Key	A string of up to 16 characters to be hashed with the cryptographic MD5 hash function.

Reverse path route caching

By default, reverse path route caching is enabled. FortiADC caches a reverse path route for inbound traffic so it can forward reply packets to the ISP link that forwarded the corresponding request packet. This is useful when your site receives traffic from multiple ISP links. For example, in [Figure 94](#), the reverse path pointer ensures that client traffic received from ISP1 is returned through ISP1.

Figure 94: Reverse path route caching enabled

When reverse path caching is not enabled, the system forwards reply packets based on the results of routing lookup.

To enable/disable reverse path route caching, use the `config router setting` CLI command:

```
FortiADC-VM # config router setting
FortiADC-VM (setting) # get
rt-cache-strict : disable
rt-cache-reverse : enable
ip-forward : enable
ip6-forward : enable
FortiADC-VM (setting) # set rt-cache-reverse disable
FortiADC-VM (setting) # end
FortiADC-VM # get router setting
rt-cache-strict : disable
rt-cache-reverse : disable
ip-forward : enable
ip6-forward : enable
```

The `rt-cache-strict` option is disabled by default. Enable it when you want to send reply packets only via the same interface that received the request packets. When enabled, source interface becomes part of the matching tuple that FortiADC uses to identify sessions, so reply traffic is forwarded from the same interface that received the traffic. (Normally each session is identified by a 5-tuple: source IP, destination IP, protocol, source port, and destination port.)

If the `rt-cache-reverse` option is enabled, you can use the `config rt-cache-reverse-exception` command to maintain an exceptions list for source IP addresses that should be handled differently. For example, if you configure an exception for 192.168.1.0/24, FortiADC will not maintain a pointer to the ISP for traffic from source 192.168.1.18. Reply packets will be forwarded based on the results of routing lookup.

```
FortiADC-docs # config router setting
FortiADC-docs (setting) # get
rt-cache-strict : disable
rt-cache-reverse : enable
ip-forward : enable
ip6-forward : enable
icmp-redirect-send : disable
FortiADC-docs (setting) # config rt-cache-reverse-exception
FortiADC-docs (rt-cache-reverse) # edit 1
Add new entry '1' for node 3740
FortiADC-docs (1) # set ip-netmask 192.168.1.0/24
FortiADC-docs (1) # end
FortiADC-docs (setting) # end
```

Packet capture

The `tcpdump` utility is supported through the CLI and web UI.

See the [FortiADC CLI Reference](#) for information on using the CLI command.

Use the following procedure to use the web UI version.

Before you begin:

- You must have a good understanding of tcpdump and filter expressions. See <http://www.tcpdump.org/manpages/pcap-filter.7.html>.
- You must have Read-Write permission for System settings.

To use the web UI version of tcpdump:

1. Go to Networking > Packet Capture.
2. Click **Add** to display an editor to specify a filter expression and other arguments.
3. Use the controls to start, stop, and download the packet capture.

Figure 95: Packet capture configuration page

Packet Capture

Interface:

IPv6: ☐ Enable

Host IP/Netmask:
Example: 192.0.2.5/24 2001:0db8:85a3::8a2e:0370:7334/64

Port:

Protocol Flag: ☒ Enable

Protocol: ☐ arp ☒ tcp ☐ udp ☐ icmp

Maximum Packet Count:

Save **Cancel**

Figure 96: Packet capture toolbar

+ Add		X Delete		Total: 1		Refresh
ID	Interface	Filter Criteria	Maximum Packet Count			
1	port1	host=172.30.144.0/22 & port=80 & protocol=tcp	10			

Chapter 18: Best Practices and Fine Tuning

This chapter is a collection of best practice tips and fine-tuning guidelines. It includes the following topics:

- [Regular backups](#)
- [Security](#)
- [Performance tips](#)
- [High availability](#)

Regular backups

Make a backup before executing disruptive operations, such as:

- Upgrading the firmware
- Running the CLI commands `execute factoryreset` or `execute restore`
- Clicking the **Reset** button in the System Information widget on the dashboard

Always password-encrypt your backups.

Security

This section lists tips to further enhance security.

Topology

- Virtual servers can be on the same subnet as physical servers. This configuration creates a one-arm load balancer.

For example, the virtual server 10.0.0.2/24 could forward to the physical server 10.0.0.3-200.

If you are deploying gradually, you might want to initially install your FortiADC in a one-arm topology during the transition phase, and route traffic to it only after you have configured FortiADC to handle it.

Long term, this is *not* recommended. Unless your network's routing configuration prevents it, it could allow clients that are aware of the physical server's IP address to bypass the FortiADC appliance by accessing the physical server directly.

- Make sure web traffic cannot bypass the FortiADC appliance in a complex network environment.
- FortiADC appliances are *not* general-purpose firewalls. While they are security-hardened network appliances, security is not their primary purpose, and you should not allow traffic to pass through without inspection. FortiADC and FortiGate complement each other to improve security, availability, and performance. To protect your servers, install the FortiADC appliance or appliances between the servers and a general purpose firewall such as a FortiGate. *FortiADC complements, and does not replace, general purpose firewalls.*
- Disable all network interfaces that should not receive any traffic.

For example, if administrative access is typically through port1, the Internet is connected to port2, and servers are connected to port3, you would disable ("bring down") port4. This would prevent an attacker with physical access from connecting a cable to port4 and thereby gaining access if the configuration inadvertently allows it.

Administrator access

- As soon as possible during initial setup, give the default administrator, `admin`, a password. This super-administrator account has the highest level of permissions possible, and access to it should be limited to as few people as possible.
- Change all administrator passwords regularly. Set a policy—such as every 60 days—and follow it. (Mark the **Change Password** check box to reveal the password dialog.)
- Instead of allowing administrative access from any source, restrict it to trusted internal hosts. On those computers that you have designated for management, apply strict patch and security policies. Always password-encrypt any configuration backup that you download to those computers to mitigate the information that attackers can gain from any potential compromise.
- Do not use the default administrator access profile for all new administrators. Create one or more access profiles with limited permissions tailored to the responsibilities of the new administrator accounts.
- By default, an administrator login that is idle for more than 30 minutes times out. You can change this to a longer period in Timeout, but Fortinet does not recommend it. Left unattended, a web UI or CLI session could allow anyone with physical access to your computer to change system settings. Small idle timeouts mitigate this risk.
- Administrator passwords should be at least 8 characters long and include both numbers and letters.
- Restrict administrative access to a single network interface (usually port1), and allow only the management access protocols needed.
- Use only the most secure protocols. Disable ping, except during troubleshooting. Disable HTTP, SNMP, and Telnet unless the network interface only connects to a trusted, private administrative network.

- Disable all network interfaces that should not receive any traffic.
- For example, if administrative access is typically through port1, the Internet is connected to port2, and servers are connected to port3, you would disable (“bring down”) port4. This would prevent an attacker with physical access from connecting a cable to port4 and thereby gaining access if the configuration inadvertently allows it.
- Immediately revoke certificates that have been compromised. If possible, automate the distribution of certificate revocation lists.

Performance tips

When configuring the system and its features, there are many settings and practices that can yield better performance.

System performance

- Delete or disable unused policies. The system allocates memory with each server policy, regardless of whether it is actually in active use. Configuring extra policies will unnecessarily consume memory and decrease performance.
- To reduce latency associated with DNS queries, use a DNS server on your local network as your primary DNS.
- If your network’s devices support them, you can create one or more VLAN interfaces. VLANs reduce the size of a broadcast domain and the amount of broadcast traffic received by network hosts, thus improving network performance.
- If you have enabled the server health check feature and one of the servers is down for an extended period, you can improve system performance by disabling group membership for the physical server, rather than allowing the server health check to continue checking for the server’s responsiveness.

Reducing the impact of logging on performance

- If you have a FortiAnalyzer, store FortiADC logs on the FortiAnalyzer to avoid resource usage associated with writing logs to the local hard disk.
- If you do not need a traffic log, disable it to reduce the use of system resources.
- Reduce repetitive log messages. Use the alert email settings to define the interval that emails are sent if the same condition persists following the initial occurrence.
- Avoid recording log messages using low severity thresholds, such as information or notification, to the local hard disk for an extended period of time. Excessive logging frequency saps system resources and can cause undue wear on the hard disk and may cause premature failure.

Reducing the impact of reports on system performance

Generating reports can be resource intensive. To avoid performance impacts, consider scheduling report generation during times with low traffic volume, such as at night and on weekends.

Keep in mind that most reports are based upon log messages. All caveats regarding log performance also apply.

Reducing the impact of packet capture on system performance

Packet capture can be useful for troubleshooting but can be resource intensive. To minimize the impact on system performance, use packet capture only during periods of minimal traffic. Use a local console CLI connection rather than a Telnet or SSH CLI connection, and be sure to stop the command when you are finished.

High availability

We recommend that you deploy high availability (HA). Keep these points in mind when setting up a cluster:

- Isolate HA interface connections from your overall network.
Heartbeat and synchronization packets contain sensitive configuration information and can consume considerable network bandwidth. For best results, directly connect the two HA interfaces using a crossover cable. If your system uses switches instead of crossover cables to connect the HA heartbeat interfaces, those interfaces must be reachable by Layer 2 multicast.
- When configuring an HA pair, pay close attention to the options ARP Packet Numbers and ARP Packet Interval. The FortiADC appliance broadcasts ARP packets to the network to ensure timely failover. Delayed broadcast intervals can slow performance. Set the value of ARP Packet Numbers no higher than needed.

When the FortiADC appliance broadcasts ARP packets, it does so at regular intervals. For performance reasons, set the value for ARP Packet Interval no greater than required.

Some experimentation might be needed to set these options at their optimum value.

We recommend that you configure an SNMP community and enable the **HA heartbeat failed** option to generate a message if the HA heartbeat fails.

Chapter 19: Troubleshooting

This chapter includes the following topics:

- [Logs](#)
- [Tools](#)
- [Solutions by issue type](#)
- [Resetting the configuration](#)
- [Restoring firmware \(“clean install”\)](#)
- [Additional resources](#)

Logs

Log messages often contain clues that can aid you in determining the cause of a problem.

Depending on the type, log messages may appear in either the event, attack, or traffic logs. The FortiADC appliance must be enabled to record event, attack, and traffic log messages; otherwise, you cannot analyze the log messages for events of that type. To enable logging of different types of events, go to Log & Report > Log Settings.

During troubleshooting, you may find it useful to lower the logging severity threshold for more verbose logs, to include more information on less severe events. To configure the severity threshold, go to Log & Report > Log Settings.

Tools

This section gives an overview of the following troubleshooting tools:

- [execute commands](#)
- [diagnose commands](#)
- [System dump](#)
- [Packet capture](#)
- [Diff](#)

execute commands

execute commands

You can use the command-line interface (CLI) execute commands to run system management utilities, such as backups, upgrades and reboots; and network diagnostic utilities, such as nslookup, ping, traceroute, and tcpdump.

The following example shows the list of execute commands:

```
FortiADC-VM # execute ?
backup backup
```

```

caching caching management
certificate certificate
checklogdisk find and auto correct errors on the log disk
clean clean
config-sync config sync
date set/get date and time
discovery-glb-virtual-server Sync virtual servers from glb server, add them to the virtual
    server list
dumpsystem dump system information for debugging purpose
dumpsystem-file manipulate the dumped debugging information
factoryreset reset to factory default
fixlogdisk correct errors on the log disk
formatlogdisk format log disk to enhance performance
geolookup lookup geography information for IP address
glb-dprox-lookup lookup GLB dynamic proximity information
glb-persistence-lookup lookup GLB persistence information
ha ha
isplookup lookup ISP name and isp-address for IP address
log log management
nslookup nslookup
packet-capture packet-capture <Port Number> [filter] (Only IPv4)
packet-capture-file packet-capture-file
packet-capture6 packet-capture6 <Port Number> [filter] (Include IPv6)
ping ping <host name | host ip>
ping-option ping option settings
ping6 ping <host name | host ipv6>
ping6-option ping6 option settings
reboot reboot the system
reload reload appliance
restore restore
shutdown shutdown appliance
ssh Simple SSH client.
statistics-db statistics db management
telnet Simple telnet client.
traceroute traceroute
vm vm
web-category-test Test a url find its web-category

```

For details, see the [FortiADC CLI Reference](#).

diagnose commands

You can use the CLI diagnose commands to gather diagnostic information that can be useful to Fortinet Customer Care when diagnosing any issues with your system. The commands are similar to the Linux commands used for debugging hardware, system, and IP networking issues.

The most important command for customers to know is `diagnose debug report`. This prepares a report you can give to your Fortinet support contact to assist in debugging an issue.

The following examples show the lists of diagnose commands:

```

FortiADC-VM # diagnose ?
debug debug
hardware hardware
llb llb
netlink netlink
server-load-balance server-load-balance

```

```

sniffer sniffer
system system

FortiADC-VM # diagnose debug ?
application set/get debug level for daemons
cli set/get debug level for CLI and CMDB
config-error-log read/clear config error information
crashlog crashlog
disable disable debug output
enable enable debug output
flow flow
info show debug info
kernel set/get debug level for kernel
report Report for tech support.
timestamp timestamp

FortiADC-VM # diagnose hardware get ?
deviceinfo list device status and information
ioport read data from an I/O port
pciconfig list information on PCI buses and connected devices
sysinfo list system hardware information

FortiADC-VM # diagnose netlink ?
backlog set netlink backlog length
device display network devices statistic information
interface netlink interface
ip ip
ipv6 ipv6
neighbor netlink neighbor
neighbor6 netlink neighbor for ipv6
route netlink routing table
route6 netlink routing table
tcp display tcp statistic information
udp display udp statistic information

FortiADC-VM # diagnose system ?
top show top process
vm check vm state

```

For details, see the [FortiADC CLI Reference](#).

System dump

The system includes utilities for generating system dump files that can help Fortinet support engineers analyze an issue for you. The CLI and Web UI versions have different usage:

- CLI—Used to dump kernel and user space information when the system is still responsive.
- Web UI—Used to dump kernel information when the system is deeply frozen.

The following is an example of CLI command usage:

```

FortiADC-VM # execute dumpsystem
This operation will reboot the system!
Do you want to continue? (y/n)y
Begins to dump userspace information
Begins to dump kernel information

```

```
FortiADC-VM # execute dumphsystem-file list
-rw----- 1 0 0 96719189 Mar 15 13:35 coredump-2016-03-15-13_35
-rw-r--r-- 1 0 0 16654391 Mar 15 13:34 user_coredump_2016_03_15_13_34_46.tar.bz2

FortiADC-VM # execute dumphsystem-file upload tftp coredump-2016-03-15-13_35 172.30.184.77
coredump-2016-03-15- 7% |** | 7152k 0:09:58 ETA
```

To use the web UI system dump utility:

1. Go to System > Debug.
2. Click **System Dump** to generate the file.

After the file has been generated, you are logged out. When you log back in and revisit the page, the system dump file appears in the file list.

3. Select the file and click **Export** to download the file.

Packet capture

The tcpdump utility is supported through the CLI and web UI.

See the [FortiADC CLI Reference](#) for information on using the CLI command.

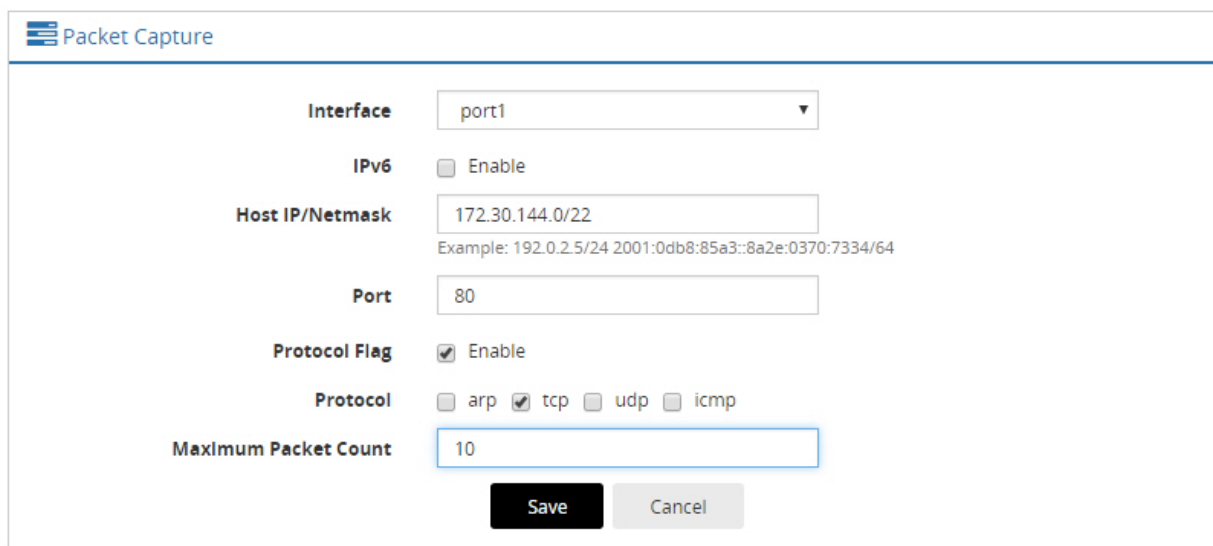
Use the following procedure to use the web UI version.

Before you begin:

- You must have a good understanding of tcpdump and filter expressions. See <http://www.tcpdump.org/manpages/pcap-filter.7.html>.
- You must have Read-Write permission for System settings.

To use the web UI version of tcpdump:

1. Go to Networking > Packet Capture.
2. Click **Add** to display an editor to specify a filter expression and other arguments.
3. Use the controls to start, stop, and download the packet capture.

Figure 97: Packet capture configuration page


The Packet Capture configuration page shows the following settings:

- Interface:** port1
- IPv6:** ☐ Enable
- Host IP/Netmask:** 172.30.144.0/22
Example: 192.0.2.5/24 2001:0db8:85a3::8a2e:0370:7334/64
- Port:** 80
- Protocol Flag:** ☒ Enable
- Protocol:** ☐ arp ☒ tcp ☐ udp ☐ icmp
- Maximum Packet Count:** 10

Buttons: Save, Cancel

Figure 98: Packet capture toolbar

+ Add		X Delete		Total: 1		Refresh
ID	Interface	Filter Criteria	Maximum Packet Count			
1	port1	host=172.30.144.0/22 & port=80 & protocol=tcp	10			

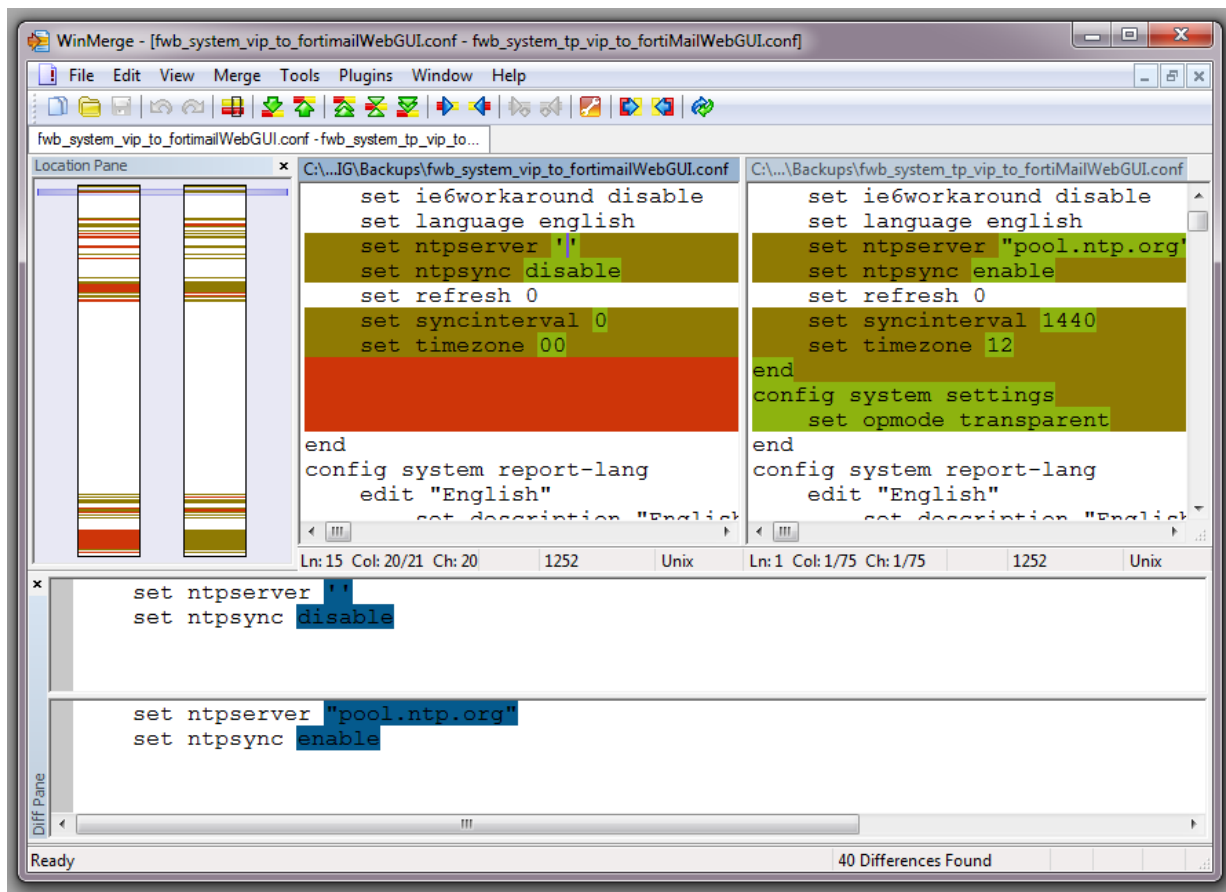
Diff

You can compare backups of the core configuration file with your current configuration. This can be useful if, for example:

A previously configured feature is no longer functioning, and you are not sure what in the configuration has changed.

You want to recreate something configured previously, but do not remember what the settings were.

Difference-finding programs, such as [WinMerge](#) and the original [diff](#) can help you to quickly find all changes. They can compare your configurations, line by line, and highlight parts that are new, modified, or deleted.

Figure 99: Configuration differences highlighted in WinMerge

For instructions, see the documentation for your diff program.

Solutions by issue type

Recommended solutions vary by the type of issue.:

- [Login issues](#)
- [Connectivity issues](#)
- [Resource issues](#)

Login issues

If an administrator is entering his or her correct account name and password, but cannot log in from some or all computers, examine that account's trusted host definitions. It should include all locations where that person is allowed to log in, such as your office, but should *not* be too broad.

Connectivity issues

One of your first tests when configuring a new policy should be to determine whether allowed traffic is flowing to your servers. Investigate the following connectivity issues if traffic does not reach the destination servers:

- Is there a FortiADC policy for the destination servers? By default, FortiADC allows traffic to reach a backend server. However, the virtual servers must also be configured before traffic can pass through.
- If your network utilizes secure connections (HTTPS) and there is no traffic flow, is there a problem with your certificate?

Checking hardware connections

If there is no traffic flowing from the FortiADC appliance, you want to rule out hardware problems.

To check hardware connections:

- Ensure the network cables are properly plugged in to the interfaces on the FortiADC appliance.
- Ensure there are connection lights for the network cables on the appliance.
- Change the cable if the cable or its connector are damaged or you are unsure about the cable's type or quality.
- Connect the FortiADC appliance to different hardware to see if that makes a difference.
- In the web UI, go to System > Networking > Interface and ensure the link status is up for the interface. If the status is down (down arrow on red circle), edit the configuration to change its status to Up.

You can also enable an interface in CLI, for example:

```
config system interface
  edit port2
    set status up
  end
```

If any of these checks solve the problem, it was a hardware connection issue. You should still perform some basic software tests to ensure complete connectivity.

If the hardware connections are correct and the appliance is powered on but you cannot connect using the CLI or web UI, you may be experiencing startup problems. See [Restoring firmware \("clean install"\)](#).

Checking routing

The `ping` and `tracert` utilities are useful for investigating issues with network connectivity and routing.

Since you typically use these tools to troubleshoot, you can allow ICMP, the protocol used by these tools, in firewall policies and on interfaces only when you need them. Otherwise, disable ICMP for improved security and performance.

By default, FortiADC appliances do not respond to `ping` and `tracert`. However, if the appliance does not respond, and there are no firewall policies that block it, ICMP type 0 (`ECHO_RESPONSE`) might be effectively disabled.

To enable ping and traceroute responses:

1. Go to Networking > Interface.
2. Select the row for the network interface and click the edit icon.
3. Under Allow Access, enable ping.
4. Save the update.

The appliance should now respond when another device such as your management computer sends a `ping` or `traceroute` to that network interface.



Note: Disabling ping only prevents the system from *receiving* ICMP type 8 (`ECHO_REQUEST`) and traceroute-related UDP. It does not disable CLI commands such as `execute ping` or `execute traceroute` that send such traffic.

To verify routes between clients and your servers:

1. Attempt to connect *through* the FortiADC appliance, from a client to a backend server, via HTTP and/or HTTPS. If the connectivity test fails, continue to the next step.
2. Use the `ping` command on both the client and the server to verify that a route exists between the two. Test traffic movement in both directions: from the client to the server, and the server to the client. Servers do not need to be able to initiate a connection, but must be able to send reply traffic along a return path. If the routing test *succeeds*, continue with step "Solutions by issue type" on page 466.

If the routing test *fails*, continue to the next step.

3. Use the `tracert` or `traceroute` command on both the client and the server (depending on their operating systems) to locate the point of failure along the route. If the route is broken when it reaches the FortiADC appliance, first examine its network interfaces and routes. To display network interface addresses and subnets, enter the CLI command:

```
show system interface
```

To display all recently-used routes with their priorities, enter the CLI command:

```
diagnose netlink route list
```

You may need to verify that the physical cabling is reliable and not loose or broken, that there are no IP address or MAC address conflicts or blacklisting, misconfigured DNS records, and otherwise rule out problems at the physical, network, and transport layer.

If these tests *succeed*, a route exists, but you cannot connect using HTTP or HTTPS, an application-layer problem is preventing connectivity.

4. For application-layer problems, on the FortiADC, examine the:
 - virtual server policy and all components it references
 - certificates (if connecting via HTTPS)
 - server service/daemon

On routers and firewalls between the host and the FortiADC appliance, verify that they permit HTTP and/or HTTPS connectivity between them.

Testing for connectivity with ping

The `ping` command sends a small data packet to the destination and waits for a response. The response has a timer that may expire, indicating that the destination is unreachable via ICMP.

ICMP is part of Layer 3 on the OSI Networking Model. `ping` sends Internet Control Message Protocol (ICMP) `ECHO_REQUEST` ("ping") packets to the destination, and listens for `ECHO_RESPONSE` ("pong") packets in reply.

Some networks block ICMP packets because they can be used in a ping flood or denial of service (DoS) attack if the network does not have anti-DoS capabilities, or because `ping` can be used by an attacker to find potential targets on the network.

Beyond basic existence of a possible route between the source and destination, `ping` tells you the amount of packet loss (if any), how long it takes the packet to make the round trip (latency), and the variation in that time from packet to packet (jitter).

If `ping` shows *some* packet loss, investigate:

- cabling to eliminate loose connections
- ECMP, split horizon, or network loops
- all equipment between the ICMP source and destination to minimize hops

If `ping` shows *total* packet loss, investigate:

- cabling to eliminate incorrect connections
- all firewalls, routers, and other devices between the two locations to verify correct IP addresses, routes, MAC lists, and policy configurations

If `ping` finds an outage between two points, use `tracert` to locate exactly where the problem is.

To use ping:

Log into the CLI via either SSH, Telnet, or the CLI Console widget of the web UI.

1. If you want to adjust the behavior of `execute ping`, first use the `execute ping-options` command.
2. Enter the command:

```
execute ping <destination_ipv4>
```

where `<destination_ipv4>` is the IP address of the device that you want to verify that the appliance can connect to, such as `192.168.1.1`.

3. If the appliance can reach the host via ICMP, output similar to the following appears:

```
PING 192.168.1.1 (192.168.1.1): 56 data bytes
64 bytes from 192.168.1.1: icmp_seq=0 ttl=253 time=6.5 ms
64 bytes from 192.168.1.1: icmp_seq=1 ttl=253 time=7.4 ms
64 bytes from 192.168.1.1: icmp_seq=2 ttl=253 time=6.0 ms
64 bytes from 192.168.1.1: icmp_seq=3 ttl=253 time=5.5 ms
64 bytes from 192.168.1.1: icmp_seq=4 ttl=253 time=7.3 ms

--- 192.168.1.1 ping statistics ---
5 packets transmitted, 5 packets received, 0% packet loss
round-trip min/avg/max = 5.5/6.5/7.4 ms
```

If the appliance cannot reach the host via ICMP, output similar to the following appears:

```
PING 10.0.0.1 (10.0.0.1): 56 data bytes
Timeout ...
Timeout ...
Timeout ...
Timeout ...
Timeout ...

--- 10.0.0.1 ping statistics ---
5 packets transmitted, 0 packets received, 100% packet loss
```

“100% packet loss” and “Timeout” indicates that the host is not reachable.



To verify that routing is bidirectionally symmetric, you should also ping the appliance.

Testing routes and latency with traceroute

The traceroute utility sends ICMP packets to test each hop along the route. It sends three packets to the destination, and then increases the time to live (TTL) setting by one, and sends another three packets to the destination. As the TTL increases, packets go one hop farther along the route until they reach the destination.

Most traceroute commands display their maximum hop count—that is, the maximum number of steps it will take before declaring the destination unreachable—before they start tracing the route. The TTL setting may result in routers or firewalls along the route timing out due to high latency.

Where ping only tells you if the signal reached its destination and returned successfully, traceroute shows each step of its journey to its destination and how long each step takes. If you specify the destination using a domain name, the traceroute output can also indicate DNS problems, such as an inability to connect to a DNS server.

By default, the traceroute utility uses UDP with destination ports numbered from 33434 to 33534. The traceroute utility usually has an option to specify use of ICMP ECHO_REQUEST (type 8) instead, as used by the Windows tracert utility. If you have a firewall and you want traceroute to work from both machines (Unix-like systems and Windows) you will need to allow both protocols inbound through your firewall (UDP ports 33434 - 33534 and ICMP type 8).

To use traceroute:

1. Log into the CLI via either SSH, Telnet, or the CLI Console widget of the web UI.
2. Enter the command:

```
execute traceroute {<destination_ipv4> | <destination_fqdn>}
where {<destination_ipv4> | <destination_fqdn>} is a choice of either the device's IP
address or its fully qualified domain name (FQDN).
```

For example, you might enter:

```
execute traceroute www.example.com
```

If the appliance *has* a complete route to the destination, output similar to the following appears:

```
traceroute to www.fortinet.com (66.171.121.34), 32 hops max, 84 byte packets
 1 172.16.1.2 0 ms 0 ms 0 ms
 2 209.87.254.221 <static-209-87-254-221.storm.ca> 2 ms 2 ms 2 ms
 3 209.87.239.129 <core-2-g0-1-1104.storm.ca> 2 ms 1 ms 2 ms
 4 67.69.228.161 2 ms 2 ms 3 ms
 5 64.230.164.17 <core2-ottawa23_POS13-1-0.net.bell.ca> 3 ms 3 ms 2 ms
 6 64.230.132.234 <core2-ottawac_POS5-0-0.net.bell.ca> 20 ms 20 ms 20 ms
 7 64.230.132.58 <core4-toronto21_POS0-12-4-0.net.bell.ca> 24 ms 21 ms 24 ms
 8 64.230.138.154 <bx4-toronto63_so-2-0-0-0.net.bell.ca> 8 ms 9 ms 8 ms
 9 64.230.185.145 <bx2-ashburn_so2-0-0.net.bell.ca> 23 ms 23 ms 23 ms
10 12.89.71.9 23 ms 22 ms 22 ms
11 12.122.134.238 <cr2.wswdc.ip.att.net> 100 ms 12.123.10.130 <cr2.wswdc.ip.att.net>
   101 ms 102 ms
12 12.122.18.21 <cr1.cgcil.ip.att.net> 101 ms 100 ms 99 ms
13 12.122.4.121 <cr1.sffca.ip.att.net> 100 ms 98 ms 100 ms
14 12.122.1.118 <cr81.sj2ca.ip.att.net> 98 ms 98 ms 100 ms
15 12.122.110.105 <gar2.sj2ca.ip.att.net> 96 ms 96 ms 96 ms
```

```

16 12.116.52.42 94 ms 94 ms 94 ms
17 203.78.181.10 88 ms 87 ms 87 ms
18 203.78.181.130 90 ms 89 ms 90 ms
19 66.171.121.34 <fortinet.com> 91 ms 89 ms 91 ms
20 66.171.121.34 <fortinet.com> 91 ms 91 ms 89 ms

```

Each line lists the routing hop number, the IP address and FQDN (if any) of that hop, and the 3 response times from that hop. Typically a value of <1ms indicates a local router.

If the appliance *does not* have a complete route to the destination, output similar to the following appears:

```

traceroute to 10.0.0.1 (10.0.0.1), 32 hops max, 84 byte packets
1 172.16.1.2 0 ms 0 ms 0 ms
2 172.16.1.10 0 ms 0 ms 0 ms
3 * * *
4 * * *

```

The asterisks (*) indicate no response from that hop in the network routing.

Examining the routing table

When a route does not exist, or when hops have high latency, examine the routing table. The routing table is where the FortiADC appliance caches recently used routes.

If a route is cached in the routing table, it saves time and resources that would otherwise be required for a route lookup. If the routing table is full and a new route must be added, the oldest, least-used route is deleted to make room.

To check the routing table in the CLI, enter:

```
diagnose netlink route list
```

Examining server daemons

If a route exists, but you cannot connect to the web UI using HTTP or HTTPS, an application-layer problem is preventing connectivity.

Verify that you have enabled HTTPS and/or HTTP on the network interface. Also examine routers and firewalls between the host and the FortiADC appliance to verify that they permit HTTP and/or HTTPS connectivity between them. Finally, you can also use the CLI command to verify that the daemons for the web UI and CLI, such as `sshd`, `cli`, `nginx`, and `php-fpm` are running and not overburdened:

```
diagnose system top delay 10
```

Checking port assignments

If you are attempting to connect to FortiADC on a given network port, and the connection is expected to occur on a different port number, the attempt will fail. For a list of ports used by FortiADC, see [Appendix B: Port Numbers](#).

Performing a packet trace

When troubleshooting malformed packet or protocol errors, it helps to look inside the protocol headers of packets to determine if they are traveling along the route you expect, and with the flags and other options you expect.



If you configure virtual servers on your FortiADC appliance, packets' destination IP addresses will be those IP addresses, not the physical IP addresses (i.e., the IP address of port1, etc.). An ARP update is sent out when a virtual IP address is configured.

If the packet trace shows that packets *are* arriving at your FortiADC appliance's interfaces but no HTTP/HTTPS packets egress, check that:

- Physical links are firmly connected, with no loose wires
- Network interfaces are brought up
- Link aggregation peers, if any, are up
- VLAN IDs, if any, match
- Virtual servers exist, and are enabled
- Matching policies exist, and are enabled
- If using HTTPS, valid server/CA certificates exist
- IP-layer and HTTP-layer routes, if necessary, match
- Servers are responsive, if server health checks are configured and enabled

Checking the SSL/TLS handshake & encryption

If the client is attempting to make an HTTPS connection, but the attempt fails after the connection has been initiated, during negotiation, the problem may be with SSL/TLS. Symptoms may include error messages such as:

- `ssl_error_no_cypher_overlap`
(Mozilla Firefox 9.0.1)
- `Error 113 (net::ERROR_SSL_VERSION_OR_CIPHER_MISMATCH): Unknown error.`
(Google Chrome 16.0.912.75 m)

The handshake is between the client and FortiADC. If the connection cannot be established, verify that the browser supports one of the key exchanges, encryption algorithms, and authentication (hashes) offered by FortiADC.

If you are not sure which cipher suites are currently supported, you can use SSL tools such as [OpenSSL](#) to discover support. For example, you could use this client-side command to know whether the server or FortiADC supports strong (HIGH) encryption:

```
openssl s_client -connect example.com:443 -cipher HIGH
```

or supports deprecated or old versions such as SSL 2.0:

```
openssl s_client -ssl2 -connect example.com:443
```

Resource issues

This section includes troubleshooting questions related to sluggish or stalled performance.

Monitoring traffic load

Heavy traffic loads can cause sustained high CPU or RAM usage. If this is unusual, no action is required. However, sustained heavy traffic load might indicate that you need a more powerful FortiADC model.

In the web UI, you can view traffic load two ways:

- Monitor current HTTP traffic on the dashboard. Go to System > Dashboard > Virtual Server and examine the throughput graphs.
- Examine traffic history in the traffic log. Go to Logs & Report > Log Browsing > Traffic Log.

DoS attacks

A prolonged denial of service (DoS) can bring your servers down if your FortiADC appliance and your network devices are not configured to prevent it. To prevent DoS attacks, enable the DoS and connection limit features. Also, configure protections on your FortiGate and other network devices. DoS attacks can use a variety of mechanisms. For in-depth protection against a wide variety of DoS attacks, you can use a specialized appliance such as FortiDDoS.

In the web UI, you can watch for attacks in two ways:

- Monitor current traffic on the dashboard. Go to System > Dashboard and examine the system-wide throughput.
- Examine attack history in the traffic log. Go to Logs & Report > Log Browsing > Security Log.

Resetting the configuration

If you will be selling your FortiADC appliance, or if you are not sure what part of your configuration is causing a problem, you can reset it to its default settings and erase data. (If you have not updated the firmware, this is the same as resetting to the factory default settings.)



Important: Back up the configuration before performing a factory reset.

To delete your data from the system, connect to the CLI and enter this command:

```
execute formatlogdisk
```

To reset the configuration, connect to the CLI and enter this command:

```
execute factoryreset
```

Restoring firmware (“clean install”)

Restoring (also called re-imaging) the firmware can be useful if:

- you are unable to connect to the FortiADC appliance using the web UI or the CLI
- you want to install firmware *without* preserving any existing configuration (i.e. a “*clean install*”)
- a firmware version that you want to install requires a different size of system partition (see the Release Notes accompanying the firmware)
- a firmware version that you want to install requires that you format the boot device (see the Release Notes accompanying the firmware)

The procedure in this section applies to physical appliances. Restoring firmware re-images the boot device. Also, restoring firmware can only be done during a boot interrupt, before network connectivity is available, and therefore requires a local console connection to the CLI. It cannot be done through an SSH or Telnet connection.



Alternatively, if you cannot physically access the appliance's local console connection, connect the appliance's local console port to a terminal server to which you have network access. Once you have used a client to connect to the terminal server over the network, you will be able to use the appliance's local console through it. However, be aware that from a remote location, you may not be able to power cycle the appliance if abnormalities occur.

For virtual appliances, you can use VMware to backup and restore virtual appliance images.



Important: Back up the configuration before performing a clean install.

To restore the firmware:

1. Download the firmware file from the Fortinet Customer Service & Support website:
<https://support.fortinet.com/>
2. Connect your management computer to the FortiADC console port using a RJ-45-to-DB-9 serial cable or a null-modem cable.
3. Initiate a local console connection from your management computer to the CLI of the FortiADC appliance, and log in as the `admin` administrator, or an administrator account whose access profile contains Read-Write permissions in the Maintenance category.
4. Connect port1 of the FortiADC appliance directly or to the same subnet as a TFTP server.
5. Copy the new firmware image file to the root directory of the TFTP server.
6. If necessary, start your TFTP server. (If you do not have one, you can temporarily install and run one such as `tftpd` (Windows, Mac OS X, or Linux) on your management computer.)



TFTP is not secure, and it does not support authentication. You should run it only on trusted administrator-only networks, and never on computers directly connected to the Internet. Turn off `tftpd` off immediately after completing this procedure.

7. Verify that the TFTP server is currently running, and that the FortiADC appliance can reach the TFTP server. To use the FortiADC CLI to verify connectivity, enter the following command:

```
execute ping 192.168.1.168
```

where `192.168.1.168` is the IP address of the TFTP server.

8. Enter the following command to restart the FortiADC appliance:

```
execute reboot
```

As the FortiADC appliances starts, a series of system startup messages appear.

```
Press any key to display configuration menu.....
```

9. Immediately press a key to interrupt the system startup.



You have only 3 seconds to press a key. If you do not press a key soon enough, the FortiADC appliance reboots and you must log in and repeat the `execute reboot` command.

If you successfully interrupt the startup process, the following messages appears:

```
[G]: Get firmware image from TFTP server.
[F]: Format boot device.
[B]: Boot with backup firmware and set as default.
[Q]: Quit menu and continue to boot with default firmware.
[H]: Display this list of options.
```

Enter G,F,B,Q,or H:

Please connect TFTP server to Ethernet port "1".

10. If the firmware version requires that you first format the boot device before installing firmware, type F. Format the boot disk before continuing.
11. Type G to get the firmware image from the TFTP server.
The following message appears:

```
Enter TFTP server address [192.168.1.168]:
```

12. Type the IP address of the TFTP server and press Enter.
The following message appears:

```
Enter local address [192.168.1.188]:
```

13. Type a temporary IP address that can be used by the FortiADC appliance to connect to the TFTP server.
The following message appears:

```
Enter firmware image file name [image.out]:
```

14. Type the file name of the firmware image and press Enter.
The FortiADC appliance downloads the firmware image file from the TFTP server and displays a message similar to the following:

```
MAC:00219B8F0D94
#####
Total 28385179 bytes data downloaded.
Verifying the integrity of the firmware image..
Save as Default firmware/Backup firmware/Run image without saving:[D/B/R]?
```



If the download fails after the integrity check with the error message:

```
invalid compressed format (err=1)
but the firmware matches the integrity checksum on the Fortinet Customer
Service & Support website, try a different TFTP server.
```

15. Type D.
The FortiADC appliance downloads the firmware image file from the TFTP server. The FortiADC appliance installs the firmware and restarts. The time required varies by the size of the file and the speed of your network connection.

The FortiADC appliance reverts the configuration to default values for that version of the firmware.
16. To verify that the firmware was successfully installed, log in to the CLI and type:

```
get system status
```

The firmware version number is displayed.

17. Either reconfigure the FortiADC appliance or restore the configuration file.

Additional resources

Fortinet also provides these resources:

- The Release Notes provided with your firmware
- [Technical documentation](#) (reference guides, installation guides, and other documents)
- [Knowledge base](#) (technical support articles)
- [Forums](#)
- [Online campus](#) (tutorials and training materials)

If you have problem using FortiADC, check within your organization first. You can save time and effort during the troubleshooting process by checking if other FortiADC administrators have experienced a similar problem before.

If you cannot resolve the issue on your own, contact [Fortinet Customer Service & Support](#).

Chapter 20: System Dashboard

The system dashboard is displayed when you log into the system (or into a virtual domain). It enables you to monitor system-wide health and utilization. You can also use it to perform some operational tasks.

This chapter discusses what you can see and do on each of the pages. It covers the following topics:

- [Status](#)
- [Data Analytics](#)
- [Server Load Balance](#)
- [Link Load Balance](#)
- [Global Load Balance](#)
- [HA](#)
- [Session Monitoring](#)

Figure 100 shows the system dashboard. Table 174 describes the information and utilities present in system dashboard portlets.

Table 174: System dashboard portlets

Portlet	Information and Utilities
Status	
System Information	<ul style="list-style-type: none"> • Hostname, current time, system uptime, serial number, firmware version. • Operations: Update firmware, upload license, reboot, shutdown, reset.
System Resources	CPU utilization, Memory utilization, disk utilization, concurrent connections, connections per second, inbound throughput, outbound throughput.
License Information	<ul style="list-style-type: none"> • License status, support contract information, and FortiGuard services version information. • Operations: Upload license, navigate to the support site, or navigate to the FortiGuard services configuration page.
Recent Event Logs	Tail of the event log.
Throughput (graph)	Shows inbound and outbound traffic.
Connections (graph)	Shows concurrent connections and connections per second.
Data Analytics	
Throughput Total	System totals for inbound and outbound throughput over the past day, hour, month, week, or 10 minutes.
Session Total	System totals for concurrent connections and connections per second over the past day, hour, month, week, or 10 minutes.

Portlet	Information and Utilities
Top Domain (Test)	Shows the top domains within the selected time frame.
Top URL	Shows the top URLs within the selected time frame.
Top Device	Shows the top devices within the selected time frame.
Top OS	Shows the top operating systems within the selected time frame.
Top Browser	Shows the top browsers within the selected time frame.
Fast Report (Widget)	Fast reports that you configure are also displayed on this tab. Fast reports show "top" reports by sessions or throughput over the past day, hour, month, week, or 10 minutes.
Server Load Balance	
Select View	<p>Displays the server load balance network in three options:</p> <ul style="list-style-type: none"> • Tree View — Shows each virtual server and the real servers in the real server pool in a tree structure. You can click each component to view or edit its configuration, or delete it. • List View—Shows each virtual server as a entry in a table. • Block View—Shows each virtual server and its related real server pool with real servers in a block. <p>Each view displays the status of the virtual server and real server pool members.</p> <p>You can filter and/or search the network map. Filters are applied before search terms. For example, if you filter by l7-load-balance type, only that set of data is searched.</p>
Add Virtual Server	<ul style="list-style-type: none"> • Status of configured virtual servers. • Select the Monitor checkbox to display throughput and connections graphs. • Click the Real Server link to display status and throughput for the virtual server's real server pool.
Add Filters	Click the button to set filters to display of
Link Load Balance	

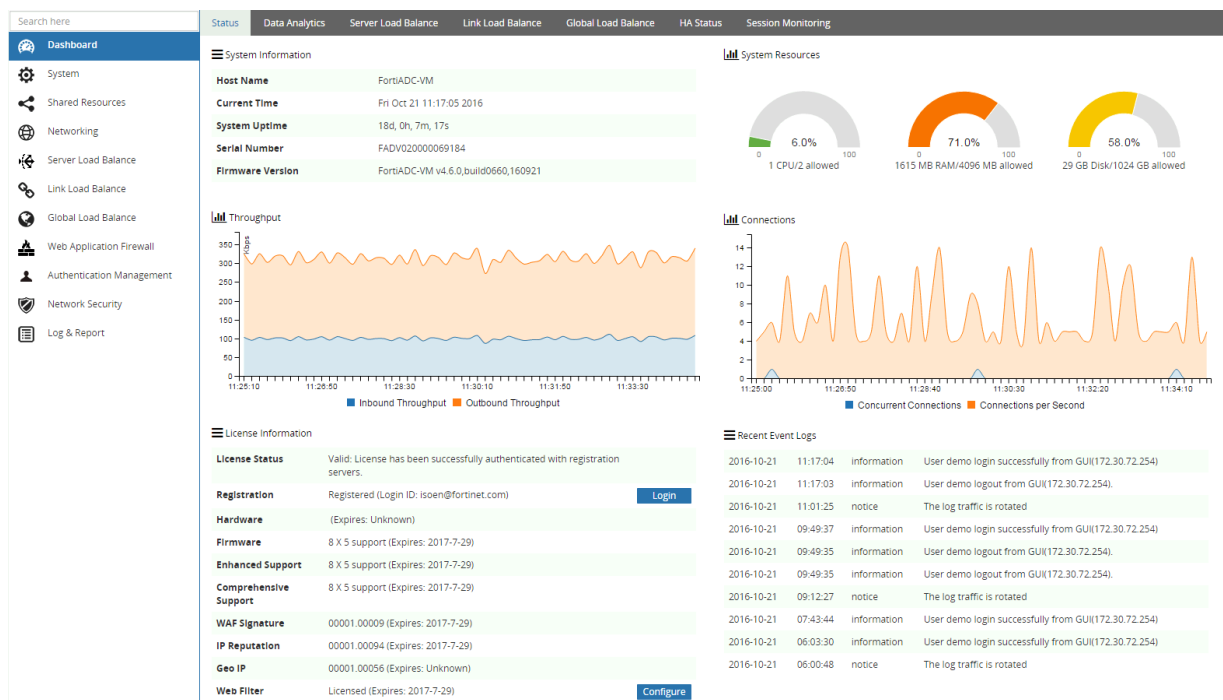
Portlet	Information and Utilities
Network Map	<p>Displays network map in three options:</p> <ul style="list-style-type: none"> • Tree View — Each virtual server is a tree. • List View—Each virtual server is a list. • Block View—Each virtual server is a block. <p>Each view displays the status of the virtual server and real server pool members.</p> <p>You can filter and/or search the network map. Filters are applied before search terms. For example, if you filter by l7-load-balance type, only that set of data is searched.</p>
Gateways	Drill-in: Select the Monitor checkbox to display throughput and session information.
Global Load Balance	
DNS responses	A table of DNS responses per server pool member.
HA	
HA Status	HA mode and details for HA member nodes.
Traffic Status	Shows information about traffic groups
Session Monitoring	
Session Table	Current sessions. You can define and apply multiple filters. After configuring filters, click OK and the table is redisplayed with matching records.
Persist Table	Current sessions. You can define and apply multiple filters. After configuring filters, click OK and the table is redisplayed with matching records.

Before you begin:

- You must have Read-Write permission for Log & Report settings.

To display the dashboard:

- Go to Dashboard.

Figure 100: System dashboard

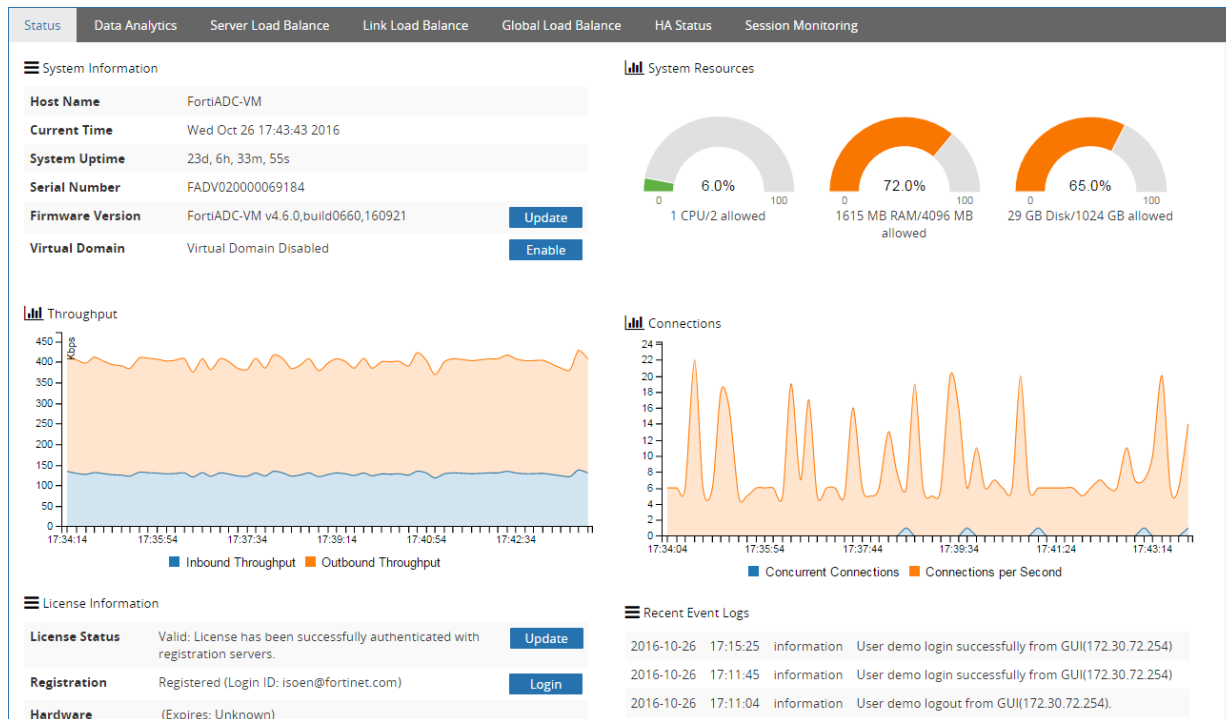
Status

By default, the Status tab opens when you select Dashboard from the side menu. The entire page is divided into six panels, each with specific information about your FortiADC system. The panels are:

- System Information
- System
- Throughput
- Connections
- License Information
- Recent Event Logs

The figure below shows a sample image of the Dashboard's Status page.

Figure 101: Status



In addition to presenting system and performance information, the Dashboard also provides following tools for managing your FortiADC:

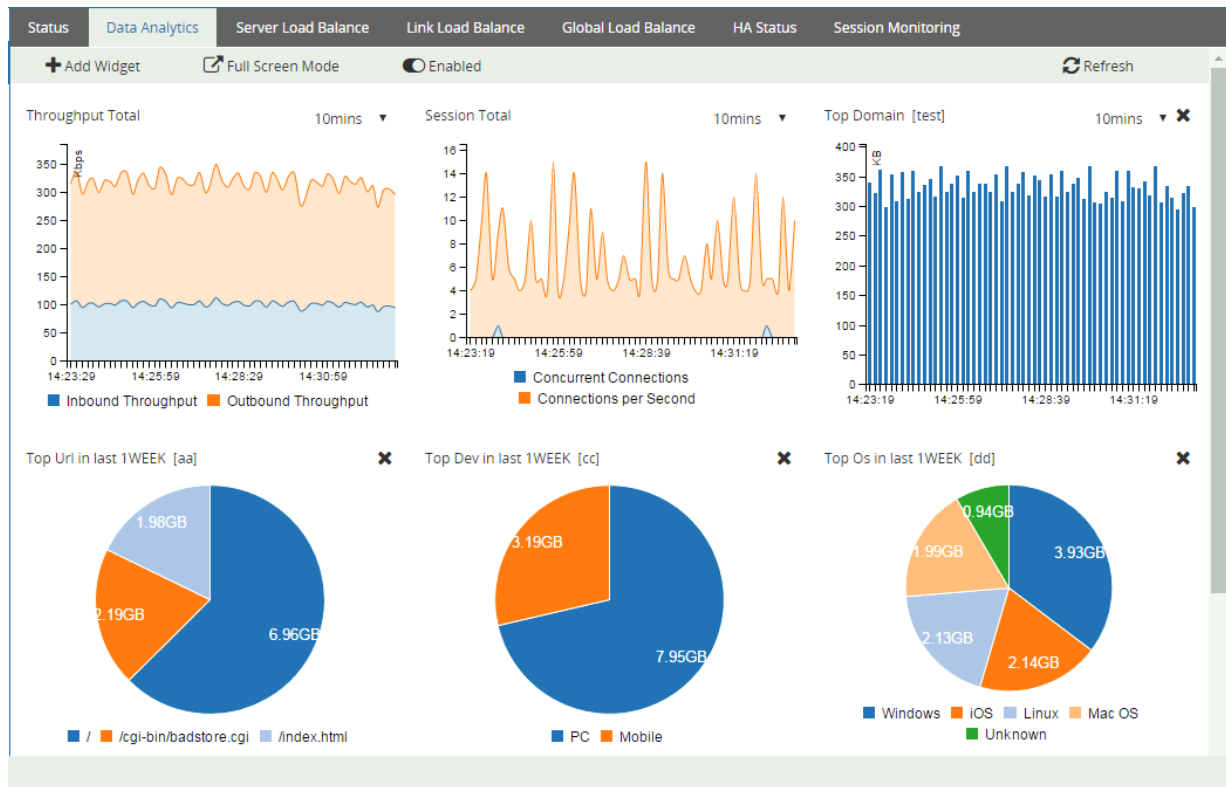
- Registration — Click the **Login** button to register your FortiADC.
- Web Filter — Click the **Configure** button to Web filters.
- Firmware — Click the **Update** button to update your FortiADC's firmware.
- Virtual Domain — Click **Enable** to enable virtual domain support and add virtual machines.
- License Status — Click **Update** to update your FortiADC license.

Data Analytics

The Data Analytics page presents system performance data in charts and graphs. Across the top of the page are three buttons that allow you to customize the data displayed and the way they are displayed. The buttons are:

- Add Widget—Click this button to create a fast report with data of your choice. See [Configure fast reports](#).
- Full Screen Mode—Click this button to open a separate window to view the data in full screen.
- Enable—Enable/disable data capture.

Figure 102: Data Analytics



Server load balance

The Server Load Balance page shows information about the virtual servers in server load-balancing operations, including their name, availability, and health check status. Here's what you can do on the Server Load Balance page:

- [Select a view option](#)
- [Filter virtual servers onscreen](#)
- [Add virtual servers](#)

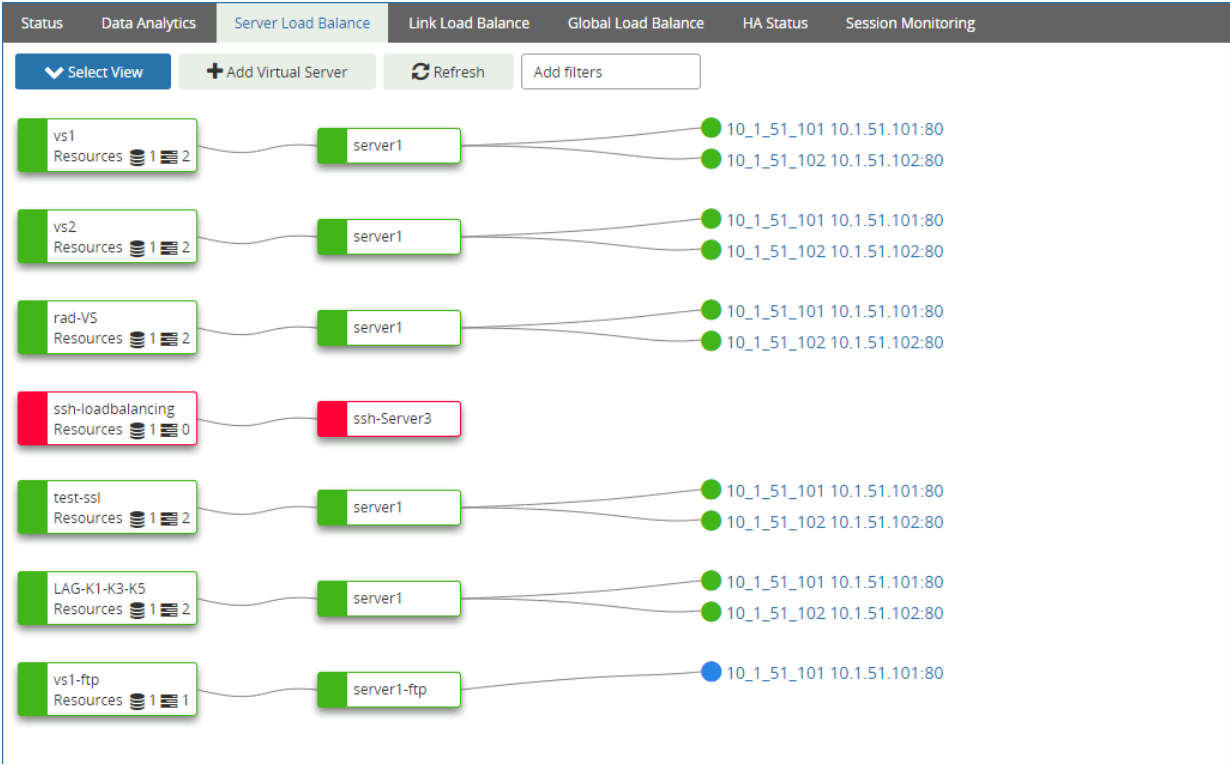
Select a display option

You can click the **Select View** button across the top of the page to choose one the following view options:

- Tree View (Default)
- List View
- Block View

The figure below shows the load-balancing servers in Tree View.

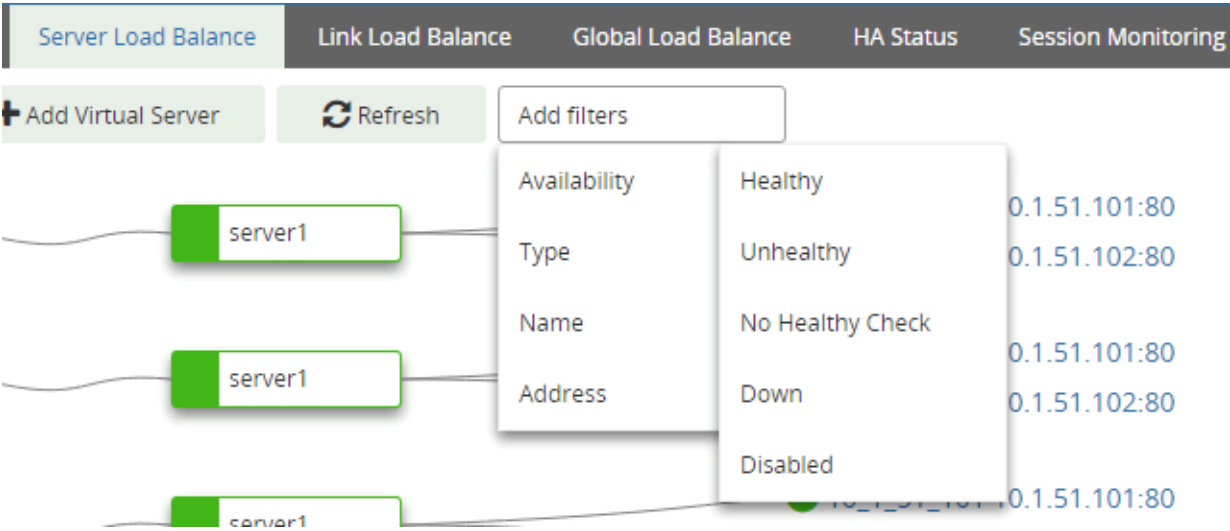
Figure 103: Server Load Balance (Tree View)



Filter virtual servers onscreen

In addition to the three view options mentioned above, you can also use the **Add Filters** button to filter the virtual servers displayed on the Server Load Balance page, as illustrated below.

Figure 104: Filter virtual servers onscreen



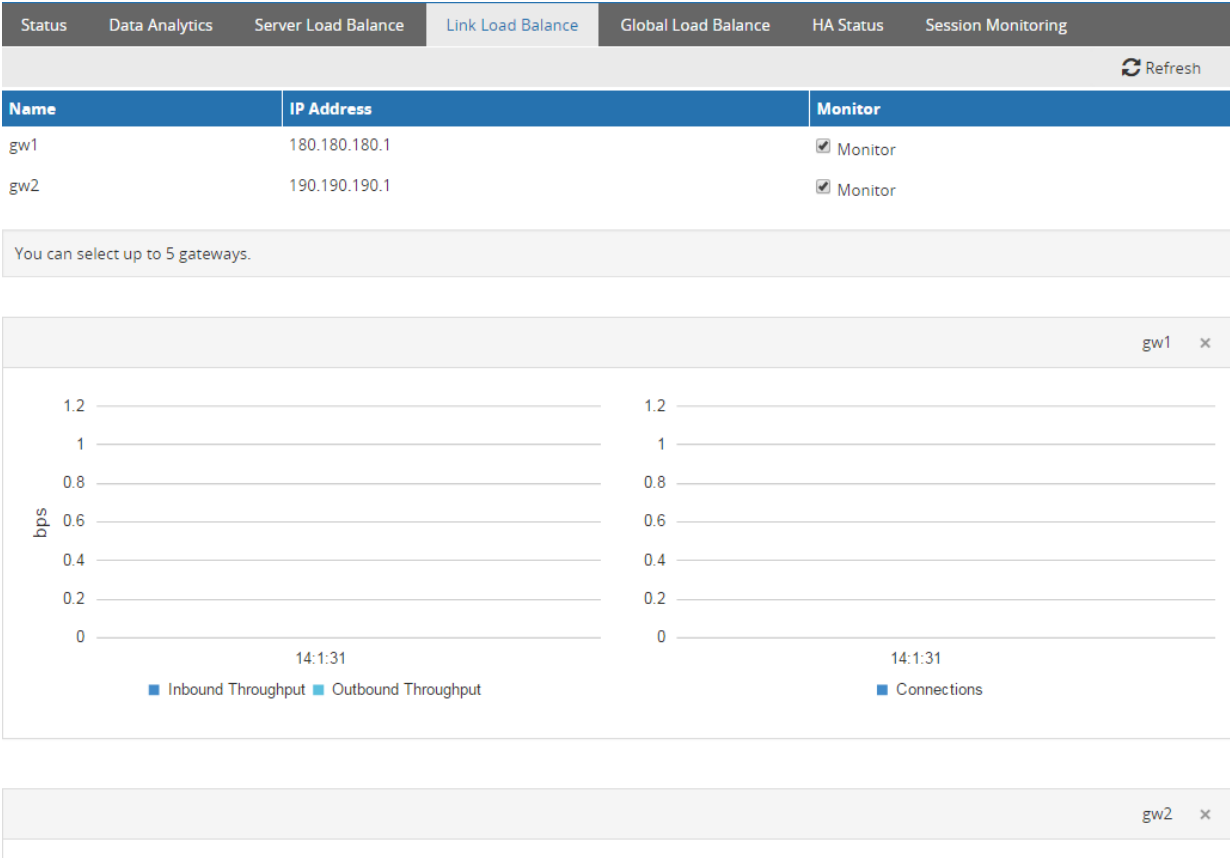
Add virtual servers

The **Add Virtual Server** button allows you to you to add virtual servers directly from the Dashboard. For instruction on how to configure virtual servers, see [Configure virtual servers](#).

Link load balance

The Link Load Balance page shows the following information about the links used in link load-balancing:


Figure 105: Link Load Balance



Global load balance

The Global Load Balance page shows the information about the servers used in global load-balancing, as illustrated below.

Figure 106: Global Load Balance

Status	Data Analytics	Server Load Balance	Link Load Balance	Global Load Balance	HA Status	Session Monitoring			
							 Refresh		
Availability	Host	Domain	VS Pool	DC Name	Server Name	VS Name	Gateway	Address	Total Response
✖	www	serform.com.	server-pool	Unknown	Unknown	Unknown	Unknown	2.3.4.10	0
✖	www	serform.com.	server-pool	Unknown	Unknown	Unknown	Unknown	::	0

HA status

The HA Status page shows the information about FortiADC's HA configuration and performance, which is divided into the following categories:

- HA Status
- Traffic Status

Figure 107: HA Status

StatusData AnalyticsServer Load BalanceLink Load BalanceGlobal Load BalanceHA StatusSession Monitoring

HA Status

Traffic Status

Mode: standalone

State: Standalone

Config Sync: N/A

Serial Number: FADV020000069184

Node ID: 0

IP Address: 169.254.88.104

Last Changed Time: Wed Dec 31 16:00:00 1969

Last Changed Reason:

Sync Statistics

Sync Pkts	Sent	Received
L4 Session and Persistence Sync Pkts	0	0
L7 Persistence Sync Pkts	0	0

Device Management Errors

Duplicate Node ID: 0

Version Mismatch: 0

Figure 108: Traffic Status

Status	Data Analytics	Server Load Balance	Link Load Balance	Global Load Balance	HA Status	Session Monitoring
HA Status						Traffic Status
						Refresh
Traffic Group Name	Current Device Node	Next Device Node	Preempt	Floating IP Addresses		
default	N/A	N/A	no	port1 10.1.50.201 port1 10.1.50.203 port2 10.1.50.201 port1 192.168.3.101 port1 10.30.30.111 port2 10.1.50.202 port2 10.1.50.201		

Session monitoring

The Session Monitor page has two tables: Session Table and Persist Table.

Figure 109: Session Table

Status	Data Analytics	Server Load Balance	Link Load Balance	Global Load Balance	HA Status	Session Monitoring					
Session Table						Persist Table					
Filter Setting						Clear	Total: 15 Refresh				
Source Address: Port	VS Address: Port	Local Address: Port	Dest Address: Port	State	Protocol	Service	In Bytes	Out Bytes	Expires	VS Name	RS Name
10.1.50.101: 52037	10.1.50.202: 80	10.1.50.101: 52037	10.1.51.101: 80	TIME_WAIT	6	tcp	483	677	2	vs2	10_1_51_101
10.1.50.101: 51939	10.1.50.202: 80	10.1.50.101: 51939	10.1.51.102: 80	TIME_WAIT	6	tcp	483	677	1	vs2	10_1_51_102
10.1.50.101: 51953	10.1.50.202: 80	10.1.50.101: 51953	10.1.51.101: 80	TIME_WAIT	6	tcp	483	677	1	vs2	10_1_51_101
10.1.50.101: 51967	10.1.50.202: 80	10.1.50.101: 51967	10.1.51.102: 80	TIME_WAIT	6	tcp	483	677	1	vs2	10_1_51_102
10.1.50.101: 51883	10.1.50.202: 80	10.1.50.101: 51883	10.1.51.102: 80	TIME_WAIT	6	tcp	483	677	0	vs2	10_1_51_102
10.1.50.101: 52065	10.1.50.202: 80	10.1.50.101: 52065	10.1.51.101: 80	TIME_WAIT	6	tcp	483	677	2	vs2	10_1_51_101

Figure 110: Persist Table

Status	Data Analytics	Server Load Balance	Link Load Balance	Global Load Balance	HA Status	Session Monitoring				
Session Table		Persist Table								
Filter Setting		Clear				Total: 0	Refresh			
Source Address	Source Port	VS Address	VS Port	Local Address	Local Port	Dest Address	Dest Port	Expires	VS Name	RS Name
1/1										

Appendix A: Fortinet MIBs

Table 175 lists the management information bases (MIBs) used with FortiADC.

Table 175: FortiADC MIBs

MIB or RFC	Description
Fortinet Core MIB	This Fortinet-proprietary MIB enables your SNMP manager to query for system information and to receive traps that are common to multiple Fortinet devices.
FortiADC MIB	This Fortinet-proprietary MIB enables your SNMP manager to query for FortiADC-specific information and to receive FortiADC-specific traps.
RFC 1213 (MIB II)	The FortiADC SNMP agent supports MIB II groups, except: There is no support for the EGP group from MIB II (RFC 1213, section 3.11 and 6.10). Protocol statistics returned for MIB II groups (IP, ICMP, TCP, UDP, and so on) do not accurately capture all FortiADC traffic activity. More accurate information can be obtained from the information reported by the FortiADC MIB.
RFC 3635 (Ethernet-like MIB)	The FortiADC SNMP agent uses any of the objects in the Ethernet-like interface types specification (dot3StatsIndex).


You can download the Fortinet MIB files from the Fortinet Customer Service & Support website, <https://support.fortinet.com/>.

To view a trap or query's name, object identifier (OID), and description, open its MIB file in a plain text editor.

To communicate with the FortiADC SNMP agent, you must first compile these MIBs into your SNMP manager. If the standard MIBs used by the SNMP agent are already compiled into your SNMP manager, you do not have to compile them again. The FortiADC SNMP implementation is read-only.

All traps sent include the message, the FortiADC appliance's serial number, and hostname.

Figure 111: FortiADC MIB download


**FORTINET**
CUSTOMER SERVICE & SUPPORT

HomeAssetAssistanceDownloadFeedback

Firmware Images

Fortinet Firmware Images And Software Releases

Welcome to the Firmware Images download center for Fortinet's extensive line of security solutions.

Select Product 

FortiADC



Release NotesDownload

Image File Path

/ FortiADC/ y4.00/ MIB/

Image Folders/Files

[Up to higher level directory](#)

	Name	Size (KB)	Date Created	Date Modified		
	FORTINET-CORE-MIB.mib	13	2015-09-08 08:09:17	2015-09-08 08:09:17	HTTPS	Checksum
	FORTINET-FORTIADC-MIB.mib	15	2015-09-08 08:09:16	2015-09-08 08:09:16	HTTPS	Checksum

Appendix B: Port Numbers

Communications between the FortiADC system, clients, servers, and FortiGuard Distribution Network (FDN) require that any routers and firewalls between them permit specific protocols and port numbers.

The following tables list the default port assignments used by the FortiADC system.

Table 176: Default ports used by FortiADC for outgoing traffic

Port Number	Protocol	Purpose
N/A	ARP	HA failover of network interfaces.
N/A	ICMP	<ul style="list-style-type: none"> • Server health checks. • <code>execute ping</code> and <code>execute traceroute</code>.
25	TCP	SMTP for alert email.
53	UDP	DNS queries.
69	UDP	TFTP for backups, restoration, and firmware updates. See commands such as <code>execute backup</code> or <code>execute restore</code> .
80	TCP	Server health checks.
123	UDP	NTP synchronization.
162	UDP	SNMP traps.
389	TCP	LDAP authentication queries.
443	TCP	<ul style="list-style-type: none"> • FortiGuard polling. • Server health checks.
514	UDP	Syslog.
6055	UDP	HA heartbeat. Layer 2 multicast.
6056	UDP	HA configuration synchronization. Layer 2 multicast.

Table 177: Default ports used by FortiADC for incoming traffic (listening)

Port Number	Protocol	Purpose
N/A	ICMP	<code>ping</code> and <code>traceroute</code> responses.
22	TCP	SSH administrative CLI access.

Port Number	Protocol	Purpose
23	TCP	Telnet administrative CLI access.
53	UDP	DNS queries from clients for global load balancing and inbound link load balancing.
80	TCP	<ul style="list-style-type: none">• HTTP administrative web UI access.• Predefined HTTP service. Only occurs if the service is used by a virtual server.
161	UDP	SNMP queries.
443	TCP	<ul style="list-style-type: none">• HTTPS administrative web UI access. Only occurs if the destination address is a network interface's IP address.• Predefined HTTPS service. Only occurs if the service is used by a virtual server, and if the destination address is a virtual server.
6055	UDP	HA heartbeat. Layer 2 multicast.
6056	UDP	HA configuration synchronization. Layer 2 multicast.

Appendix C: Scripts

You can embed Lua scripts to take actions that are not supported by the built-in feature set.

This appendix provides guidance for getting started. It includes the following topics:

- [Events and actions](#)
- [Predefined commands](#)
- [Control structures](#)
- [Operators](#)
- [String library](#)
- [Examples](#)

For general information about Lua, see <http://www.lua.org/docs.html>.

Events and actions

Scripts are associated with a particular virtual server, and they are event-driven. A script is triggered when the associated virtual server receives an HTTP request or response. Then, it does the programmed action.

[Table 178](#) provides syntax, usage, and examples of the predefined commands that are useful for writing scripts.

Table 178: Script events and actions

Event/Action	Description
Event	
HTTP_REQUEST	The virtual server receives a complete HTTP request header.
HTTP_RESPONSE	The virtual server receives a complete HTTP response header.
RULE_INIT	The event is used to initialize global or static variables used within a script. It is triggered when a script is added or modified, or when the device starts up, or when the software is restarted.
Action	
in Lua mode	An action defined by a Lua script that uses predefined commands and variables to manipulate the HTTP request/response or select a content route.

Predefined commands

[Table 179](#) provides syntax, usage, and examples of the predefined commands that are useful for writing scripts.

Table 179: Predefined commands

Syntax	Usage and Example
Global	
<code>debug("msg", ...)</code>	<p>Write the message to the debug buffer. For example:</p> <pre>debug("HTTP Request method is %s.\n", HTTP:method_get ())</pre> <p>Debug strings can be written to the console when the event is triggered. This is helpful when you are testing your scripts.</p> <p>To enable debug strings to be written to the console, use the following CLI commands:</p> <pre>diagnose debug enable diagnose debug application haproxy scripting</pre>

Syntax	Usage and Example
<code>cmp_addr(addr, addr_group)</code>	<p>Used to match one IP address against a group of IP addresses. It can automatically detect IPv4 and IPv6 and can be used to compare IPv4 addresses with IPv6 addresses.</p> <p>For example:</p> <pre>cmp_addr("192.3.2.1/24", "192.3.2.0/32") cmp_addr("::ffff:192.3.2.1/120", "::ffff:192.3.2.0/128") cmp_addr("192.3.2.1/24", "::ffff:192.3.2.0/128")</pre> <p>Input format:</p> <p>For an IPv4 <code>ip_addr/[mask]</code>, the mask can be a number between 0 and 32 or a dotted format like 255.255.255.0</p> <p>For an IPv6 <code>ip_addr/[mask]</code>, the mask can be a number between 0 and 128.</p> <p>FortiADC supports address group for the second argument.</p> <pre>when RULE_INIT{ --initialize the address group here addr_group = "192.168.1.0/24" --first network address addr_group = addr_group..",::ffff:172.30.1.0/120" --second network address --so on and so forth } when HTTP_REQUEST{ client_ip=HTTP:client_addr() match_ip=cmp_addr(client_ip, addr_group) }</pre>
<code>log("fmt", ...)</code>	<p>Writes log messages into the SLB log category in the script log part. You must enable Script log and SLB sub-category under the Script log on the log setting page. For example:</p> <pre>log("This HTTP Request method is %s.\n", HTTP:method_get())</pre> <p>Note: \ and % are handled in a unique way. Special characters that the log supports are :~!@#\$%^&*()_+{}[]. If you want to print out % in the log, you must use %%; if you want to print out \, you must use \\.</p>

Syntax	Usage and Example
rand()	<p>Generates a random number. For example:</p> <pre>a = rand() debug("a=%d\n", a)</pre>
HTTP	
header_get_names()	<p>Returns a list of all the headers present in the request or response. For example:</p> <pre>--use header and value headers = HTTP:header_get_names() for k, v in pairs(headers) do debug("The value of header %s is %s.\n", k, v) end --only use the header name for name in pairs(headers) do debug("The request/response includes header %s.\n", name) end</pre>
header_get_values(header_name)	<p>Returns a list of value(s) of the HTTP header named <header_name>, with a count for each value. Note that the command returns all the values in the headers as a list if there are multiple headers with the same name. For example:</p> <pre>cookies=HTTP:header_get_values("Cookie") for k, cnt in pairs(cookies) do debug("initially include cookie %s cnt %d\n", k, v) end</pre>
header_get_value(header_name)	<p>Returns the value of the HTTP header named <header_name>.</p> <p>Returns false if the HTTP header named <header_name> does not exist. Note: The command operates on the value of the last header if there are multiple headers with the same name. For example:</p> <pre>host = HTTP:header_get_value("Host")</pre>
header_remove(header_name)	<p>Removes all headers names with the name <header_name>. For example:</p> <pre>HTTP:header_remove("Cookie")</pre>

Syntax	Usage and Example
<code>header_remove2(header_name,-countid)</code>	<p><code>header_get_values()</code> returns an count ID for each item. This count ID can be used in both <code>header_remove2()</code> and <code>header_replace2()</code> to remove and replace a certain header of a given name referenced by the count ID. For example:</p> <pre> cookies=HTTP:header_get_values("Set-Cookie") for k, v in pairs(cookies) do debug("include cookie %s cnt %d\n", k, v) end if HTTP:header_remove2("Set-Cookie", 1) then debug("remove 1st cookie\n") end </pre>
<code>header_insert(header_name, value)</code>	<p>Inserts the named HTTP header(s) and value(s) into the end of the HTTP request or response. For example:</p> <pre> HTTP:header_insert("Cookie", "cookie=server1") </pre>
<code>header_replace(header_name, value)</code>	<p>Replaces the value of the last occurrence of the header named <header_name> with the string <value>. Performs a header insertion if the header is not present. For example:</p> <pre> HTTP:header_replace("Host", "www.fortinet.com") </pre>
<code>header_replace2(header_name, value,countid)</code>	<p><code>header_get_values()</code> returns an count ID for each item. This count ID can be used in both <code>header_remove2()</code> and <code>header_replace2()</code> to remove and replace a certain header of a given name referenced by the count ID. For example:</p> <pre> cookies=HTTP:header_get_values("Set-Cookie") for k, v in pairs(cookies) do debug("include cookie %s cnt %d\n", k, v) end if HTTP:header_replace2("Set-Cookie", "new2=value2", 2) then debug("replace 2nd cookie by new2=value2\n") end </pre>
<code>header_exists(header_name)</code>	<p>Returns true if the named header is present and not empty on the request or response. For example:</p> <pre> if HTTP:header_exists("Cookie") then ... end </pre>
<code>header_count(header_name)</code>	<p>Returns the number of HTTP headers present in the request or response. For example:</p> <pre> count = HTTP:header_count("Cookie") </pre>

Syntax	Usage and Example
<code>method_get()</code>	<p>Return the string of the HTTP request method. For example:</p> <pre>method = HTTP:method_get()</pre>
<code>method_set(string)</code>	<p>Set the HTTP request method to the string "<i>value</i>". For example:</p> <pre>HTTP:method_set("POST")</pre>
<code>path_get()</code>	<p>Returns the path part of the HTTP request. For example:</p> <pre>path = HTTP:path_get()</pre>
<code>path_set(string)</code>	<p>Sets the path part of the HTTP request. The client will not see the update unless the web application uses the requested path to generate response headers and/or content. If you want the client to see the update to the path in the browser's address bar, you can send an HTTP redirect using <code>HTTP:redirect</code> or <code>HTTP:respond</code>. For example:</p> <pre>HTTP:path_set("/other.html")</pre>
<code>uri_get()</code>	<p>Returns the URI given in the request. For example:</p> <pre>uri = HTTP:uri_get()</pre>
<code>uri_set(string)</code>	<p>Changes the URI passed to the server. It should always start with a slash. For example:</p> <pre>HTTP:uri_set("/index.html?value=xxxx")</pre>
<code>query_get()</code>	<p>Returns the query part of the HTTP request. For example:</p> <pre>query = HTTP:query_get()</pre>
<code>query_set(string)</code>	<p>Sets the query part of the HTTP request. For example:</p> <pre>HTTP:query_set("value=xxx")</pre>
<code>redirect("URL", ...)</code>	<p>Redirects an HTTP request or response to the specified URL. For example:</p> <pre>Host = HTTP:header_get_value("host") Path = HTTP:path_get() HTTP:redirect("https://%s%s", Host, Path)</pre>
<code>redirect_with_cookie(URL, cookie)</code>	<p>Redirects an HTTP request or response to the specified URL with Cookie. For example:</p> <pre>HTTP:redirect_with_cookie("www.example.com", "server=nginx")</pre>

Syntax	Usage and Example
<code>version_get()</code>	Returns the HTTP version of the request or response. For example: <code>vers = HTTP:version_get()</code>
<code>version_set(string)</code>	Sets the HTTP version of the request or response. For example: <code>HTTP:version_set("1.0")</code>
<code>status_code_get()</code>	Returns the response status code output as string. For example: <code>responsestatus=HTTP:status_code_get()</code>
<code>status_code_set(string)</code>	Sets the response status code. For example: <code>HTTP:status_code_set("301")</code>
<code>code_get()</code>	Returns the response status code,output as integer. For example: <code>responsestatus=HTTP:code_get()</code>
<code>code_set(integer)</code>	Sets the response status code. For example: <code>HTTP:code_set(301)</code>
<code>reason_get()</code>	Returns the response reason. For example: <code>HTTP:reason_get()</code>
<code>reason_set(string)</code>	Sets the response reason. For example: <code>HTTP:reason_set(string)</code>
<code>rand_id()</code>	Returns a random string of 32-long in hex format, which can be inserted directly as an HTTP header. For example: <code>ID=HTTP:rand_id()</code> <code>HTTP:header_insert("Message-ID", ID)</code>
<code>client_addr()</code>	Returns the client IP address of a connection for an HTTP_REQUEST packet, which is the source address for the HTTP_REQUEST packet. It's a destination address. For example: <code>CIP=HTTP:client_addr()</code>
<code>local_addr()</code>	For HTTP_REQUEST, returns the IP address of the virtual server the client is connected to; for HTTP_RESPONSE, returns the incoming interface IP address of the return packet. For example: <code>LIP=HTTP:local_addr()</code>

Syntax	Usage and Example
<code>remote_addr()</code>	Returns the IP address of the host on the far end of the connection. For example: <code>RIP=HTTP:remote_addr()</code>
<code>server_addr()</code>	Returns the IP address of the server in HTTP_RESPONSE. <code>SIP=HTTP:server_addr()</code>
<code>close()</code>	Closes an HTTP connection using code 503. For example: <code>HTTP:close()</code>
Load Balance	
<code>routing(content_route)</code>	Selects a content route. For example: <code>LB:routing("content2")</code>

Control structures

Table 180 lists Lua control structures.

Table 180: Lua control structures

Type	Structure
if then else	<pre> if condition1 then ... elseif condition2 then ... else ... end </pre>
for	<pre> --fetch all values of table 't' for k, v in pairs(t) do ... end </pre>

Operators

Table 181 lists the FortiADC operators.

Table 181: Lua operators

FortiADC Operator	Operator sub-type	Description
- +	Arithmetic	Unary minus, unary plus.
~	Bitwise	Bitwise NOT.
not	Logical	Performs a logical "not" on a value.
* / %	Arithmetic	Multiple, divide, remainder.
//		Floor division.
^		Exponentiation.
+ -	Arithmetic	Add and subtract.
<< >>	Bitwise	Left and right shift.
< > <= >=	Relational	Boolean less, greater, less than or equal, and greater than or equal.
== !=	Relational	Boolean equal and not equal.
&	Bitwise	Bitwise AND.
~	Bitwise	Bitwise exclusive OR.
	Bitwise	Bitwise OR.
and	Logical	Performs a logical "and" comparison between two values.
or	Logical	Performs a logical "or" comparison between two values.
starts_with (a,b)	String	Tests to see if String a starts with String b.
ends_with (a,b)	String	Tests to see if String a ends with String b.

FortiADC Operator	Operator sub-type	Description
..		The string concatenation operator in Lua is denoted by two dots ('..'). If both operands are strings or numbers, then they are converted to a string. It's the same as <code>_concat</code> .

String library

The FortiADC OS supports only the [Lua string library](#). All other libraries are disabled. The string library includes the following string-manipulation functions:

- `string.byte(s, i)`
- `string.char(i1,i2...)`
- `string.dump(function)`
- `string.find(s, pattern)`
- `string.format`
- `string.gmatch`
- `string.gsub`
- `string.len`
- `string.lower`
- `string.match`
- `string.rep`
- `string.reverse`
- `string.sub`
- `string.upper`
- `string.starts_with`
- `string.ends_with`

For example: `uri:starts_with (b), uri:ends_with (b)`

Note:

- If you want to do regular expression match, you can use `string.match` with Lua patterns.
- All relational operators `>`, `<`, `>=`, `<=`, `~=`, `==` apply to strings. Especially, `==` can be used to test if one string equals to another string.
- `string.find` can be used to test whether one string contains another string.

For a tutorial on scripting with the Lua string library, see <http://lua-users.org/wiki/StringLibraryTutorial>.

Examples

This section provides example scripts for popular use cases. It includes the following examples:

- [Select content routes based on URI string matches](#)
- [Rewrite the HTTP request host header and path](#)
- [Rewrite the HTTP response Location header](#)
- [Redirect HTTP to HTTPS using Lua string substitution](#)
- [Redirect mobile users to the mobile version of a website](#)



Tip: The examples show debug strings. Debug strings can be written to the console when the event is triggered. This is helpful when you are testing your scripts.

To enable debug strings to be written to the console, use the following CLI commands:

```
diagnose debug enable
diagnose debug application haproxy scripting
```

Select content routes based on URI string matches

The content routing feature has rules that match HTTP requests to content routes based on a Boolean AND combination of match conditions. If you want to select routes based on a Boolean OR, you can configure multiple rules. The content routing rules table is consulted from top to bottom until one matches.

In some cases, it might be simpler to get the results you want using a script. In the following example, each rule selects content routes based on OR match conditions.

```
-- Content routing example
when RULE_INIT {
  debug("get header init 1\n")
}

when HTTP_REQUEST{
  uri = HTTP:uri_get()
  if uri:find("sports") or uri:find("news") or uri:find("government") then
    LB:routing("sp2")
    debug("uri %s matches sports|news|government\n", uri);
  elseif uri:find("finance") or uri:find("technology") or uri:find("shopping") then
    LB:routing("sp3")
    debug("uri %s matches finance|technology|shopping\n", uri);
  elseif uri:find("game") or uri:find("bbs") or uri:find("testing") then
    LB:routing("sp4")
    debug("uri %s matches game|bbs|testing\n", uri);
  elseif uri:find("billing") or uri:find("travel") or uri:find("weibo") then
    LB:routing("sp5")
    debug("uri %s matches billing|travel|weibo\n", uri);
  else
    debug("no matches for uri: %s \n", uri);
  end
end
}
```

To use a script for content routing:

1. Create the content route configuration objects. In the example above, sp2, sp3, sp4, and sp4 are the names of the content route configuration objects. You do not need to configure matching conditions for the content routes, however, because the script does the content matching.
2. Create a script that matches content to the content route configuration objects, as shown above. Create a configuration object for the script.
3. In the virtual server configuration:
 - a. Enable content routing and select the content route configuration objects.
 - b. Select the script.

Rewrite the HTTP request host header and path

You can use the content rewriting feature to rewrite the HTTP request Host header or the HTTP request URL. If you need more granular capabilities, you can use scripts. The following example rewrites the HTTP Host header and path.

```
-- Rewrite the HTTP Host header and path in a HTTP request
when RULE_INIT {
  debug("rewrite the HTTP Host header and path in a HTTP request \n")
}

when HTTP_REQUEST{
  host = HTTP:header_get_value("Host")
  path = HTTP:path_get()
  if host:lower():find("myold.hostname.com") then
    debug("found myold.hostname.com in Host %s \n", host)
    HTTP:header_replace("Host", "mynew.hostname.com")
    HTTP:path_set("/other.html")
  end
}
```

Note: You might find it useful to use a combination of string manipulation functions. For example, this script uses `lower()` to convert the Host strings to lowercase in combination with `find()`, which searches for the Host header for a match: `host:lower():find("myold.hostname.com")`.

Rewrite the HTTP response Location header

You can use the content rewriting feature to rewrite the HTTP response Location header. If you are more comfortable using Lua string substitution, you can write a script to get the results you want. The following example rewrites the HTTP response Location header.

```
-- Rewrite the HTTP body in the response
when RULE_INIT {
  debug("rewrite the HTTP response replacing myold.hostname.com with mynew.hostname.com \n")
}

when HTTP_RESPONSE{
  location = HTTP:header_get_value("Location")
  if location:lower():find("myold.hostname.com") then
    debug("found myold.hostname.com in Location %s \n", location)
    HTTP:header_replace("Location", "mynew.hostname.com")
  end
}
```

Redirect HTTP to HTTPS using Lua string substitution

You can use the content rewriting feature to redirect an HTTP request to an HTTPS URL that has the same host and request URL using a PCRE regular expression. If you are more comfortable using Lua string substitution, you can write a script to get the results you want. The following example redirects users to the HTTPS location.

```
-- Redirect HTTP to HTTPS
when RULE_INIT {
  debug("http to https redirect\n")
}

when HTTP_REQUEST{
  host = HTTP:header_get_value("Host")
  path = HTTP:path_get()
  HTTP:redirect("https://%s%s",host,path);
}
```

Redirect mobile users to the mobile version of a website

The content rewriting feature does not support matching the User-Agent header. You can write a script that detects User-Agent headers that identify mobile device users and redirect them to the mobile version of a website.

```
-- Redirect mobile users to the mobile version of a website by parsing the User-Agent
  header
when RULE_INIT {
  debug("detect User-Agent and go to mobile site\n")
}

when HTTP_REQUEST{
  path = HTTP:path_get()
  debug("path=%s\n",path)
  agent = HTTP:header_get_value("User-Agent")
  if agent:lower():find("iphone") or agent:lower():find("ipad") then
    debug("found iphone or ipad in User-Agent %s \n", agent)
    HTTP:redirect("https://m.mymobilesite.com%s",path)
  end
}
```

Appendix D: Maximum Configuration Values

This table shows the maximum number of configuration objects or limits that vary by them, and are not a guarantee of performance. For values such as hardware specifications that do not vary by software version or configuration, see your model's QuickStart Guide or datasheet.

Table 182: Maximum configuration objects - Hardware models

		200D/100F	300D	400D	700D	1500D	2000D	4000D
System								
Administration	Administrative users	300	300	300	300	300	300	300
	Access profiles	16	64	64	64	64	64	64
	Virtual domains (VDOMs)	10	10	10	30	45	60	90
Certificates	Any configuration object	256	256	256	256	256	256	256
Shared Resources	Address	1024	2048	2048	2048	2048	2048	4096
	Address group	256	256	256	256	256	256	256
	Health checks	128	256	256	256	256	256	512
	ISP address book	32	32	32	32	32	32	32
	Schedule	256	256	256	256	256	256	256
	Schedule group	64	64	64	64	64	64	64
	Service	1024	2048	2048	2048	2048	2048	4096
	Service group	256	256	256	256	256	256	256
SNMP	SNMP community	16	16	16	16	16	16	16
	SNMP community Host	16	16	16	16	16	16	16
	SNMP user	16	16	16	16	16	16	16
Networking								
Interface	Physical network interfaces	4	4	4	12	12	20	24
	VLAN interfaces	256	512	512	512	512	512	1024

		200D/100F	300D	400D	700D	1500D	2000D	4000D
Routing	ARP table entries (per VDOM)	4096	4096	4096	4096	4096	4096	4096
	Static routes	2048	4096	4096	4096	4096	4096	4096
	Policy routes	64	128	128	128	128	128	256
	ISP routes	32	32	32	32	32	32	32
NAT	Any configuration object	256	256	256	256	256	256	256
QoS	Any configuration object	256	256	256	256	256	256	256
Packet capture	Table	5	5	5	5	5	5	5
User								
	Any configuration object	256	256	256	256	256	256	256
Server Load Balancing								
Virtual Servers		1024	2048	2048	2048	2048	2048	4096
Real Server Pool	Pools	1024	2048	2048	2048	2048	2048	4096
	Pool members	1024	2048	2048	2048	2048	2048	4096
	Real server SSL profiles	256	256	256	256	256	256	256
Resources	Profiles	256	256	256	256	256	256	256
	Cache policies	256	256	256	256	256	256	256
	Compression policies	256	256	256	256	256	256	256
	Persistence policies	128	256	256	256	256	256	512
	Method policies	64	128	128	128	128	128	256
	Authentication policies	256	256	256	256	256	256	256
	Scripts	256	256	256	256	256	256	256
Content Rules	Content routing rules	256	512	512	512	512	512	1024
	Content rewriting rules	256	512	512	512	512	512	1024

		200D/100F	300D	400D	700D	1500D	2000D	4000D
Link Load Balancing								
Link Group	Gateway	1024	2048	2048	2048	2048	2048	4096
	Link group	512	1024	1024	1024	1024	1024	2048
	Link group member	1024	2048	2048	2048	2048	2048	4096
Virtual Tunnel Group	Virtual tunnel group	512	1024	1024	1024	1024	1024	2048
	Virtual tunnel member	256	256	256	256	256	256	256
Policy	LLB policy rule	512	1024	1024	1024	1024	1024	2048
Global Load Balancing								
Any configuration object		256	256	256	256	256	256	256
Security								
Any configuration object		256	256	256	256	256	256	256
Log & Report								
Remote Syslog Servers		3	3	3	3	3	3	3

Table 183: Maximum configuration objects - Virtual Appliances

		VM01	VM02	VM04	VM08
System					
Administration	Administrative users	300	300	300	300
	Access profiles	8	16	64	64
	Virtual domains (VDOMs)	0	0	5	10
Certificate	Any configuration object	256	256	256	256

		VM01	VM02	VM04	VM08
Shared Resources	Address	512	1024	2048	4096
	Address group	256	256	256	256
	Health checks	64	128	256	512
	ISP address book	32	32	32	32
	Schedule	256	256	256	256
	Schedule group	64	64	64	64
	Service	512	1024	2048	4096
	Service group	256	256	256	256
SNMP	SNMP community	16	16	16	16
	SNMP community host	16	16	16	16
	SNMP user	16	16	16	16
Networking					
Interfaces	Physical network interfaces	10	10	10	10
	VLAN interfaces	128	256	512	1024
Routing	ARP table entries (per VDOM)	4096	4096	4096	4096
	Static routes	1024	2048	4096	4096
	Policy routes	32	64	128	256
	ISP routes	32	32	32	32
NAT	Any configuration object	256	256	256	256
QoS	Any configuration object	256	256	256	256
Packet Capture	Table	5	5	5	5
User					
	Any configuration object	256	256	256	256
Server Load Balancing					

		VM01	VM02	VM04	VM08
Virtual Servers		512	1024	2048	4096
Real Server Pool	Pools	512	1024	2048	4096
	Pool members	512	1024	2048	4096
Resources	Real server SSL profile	256	256	256	256
	Profiles	256	256	256	256
	Cache policies	256	256	256	256
	Compression policies	256	256	256	256
	Persistence policies	128	128	128	256
	Method policies	32	64	128	256
	Authentication policies	256	256	256	256
	Scripts	256	256	256	256
Content Rules	Content routing rules	128	256	512	1024
	Content rewriting rules	128	256	512	1024
Link Load Balancing					
Link Group	Gateway	512	1024	2048	4096
	Link group	256	512	1024	2048
	Link group member	512	1024	2048	4096
Virtual Tunnel	Virtual tunnel	256	512	1024	2048
	Virtual tunnel member	256	256	256	256
Policy	LLB policy rule	256	512	1024	2048
Global Load Balancing					
Any configuration object		256	256	256	256
Security					
Any configuration object		256	256	256	256

	VM01	VM02	VM04	VM08
Log & Report				
Remote Syslog Servers	3	3	3	3

Appendix E: High Speed Logging Binary Format

The high speed logging feature sends a binary log file. It has the following structure:

```
typedef struct __high_speed_log_header {
    unsigned char    msg_ver;
    unsigned char    have_dev_vd_str;
    unsigned char    dev;
    unsigned char    timezone;
    unsigned int     vid;
    unsigned char    cmd;
    unsigned char    log_type;
    unsigned char    log_subtype;
    unsigned short   total_len;
    unsigned char    total_count;
    unsigned short   next_len;
} high_speed_log_header_t;

typedef struct __dev_vdom_str {
    unsigned short   next_len;
    unsigned char    dev_len;
    unsigned char    vdom_len;
    unsigned char    data[]; //dev+vdom
} dev_vdom_str_t;

typedef struct _tlog_l4 {
    unsigned short   next_len;
    unsigned int     itime;
    unsigned int     duration;
    uint64_t         ibytes;
    uint64_t         obytes;
    unsigned char    proto;
    unsigned char    srv;
    char             src_ver;
    unsigned short   src_port;
    char             dst_ver;
    unsigned short   dst_port;
    char             trans_src_ver;
    unsigned short   trans_src_port;
    char             trans_dst_ver;
    unsigned short   trans_dst_port;
    unsigned short   policy;
    unsigned char    policy_len;
    unsigned char    action;
    unsigned char    data[]; //src+dst+policy
} tlog_l4_t;
```

```

typedef struct _tlog_fw {
    unsigned short    next_len;
    unsigned int      itime;
    unsigned int      duration;
    uint64_t          ibytes;
    uint64_t          obytes;
    unsigned char     proto;
    unsigned char     srv;
    char              src_ver;
    unsigned short    src_port;
    char              dst_ver;
    unsigned short    dst_port;
    char              trans_src_ver;
    unsigned short    trans_src_port;
    char              trans_dst_ver;
    unsigned short    trans_dst_port;
    unsigned short    policy;
    unsigned char     policy_len;
    char              nat_policy; //reserved for dictionary
    char              nat_policy_len;
    unsigned short    llb_policy;
    unsigned char     llb_policy_len;
    unsigned short    llb_link;
    unsigned char     llb_link_len;
    unsigned char     action;
    unsigned char     data[]; /*osip:odip:rsip:rdip:nat...*/
} tlog_fw_t;

typedef struct _tlog_llb {
    unsigned short    next_len;
    unsigned int      itime;
    unsigned int      duration;
    uint64_t          ibytes;
    uint64_t          obytes;
    unsigned char     proto;
    unsigned char     srv;
    char              src_ver;
    unsigned short    src_port;
    char              dst_ver;
    unsigned short    dst_port;
    char              trans_src_ver;
    unsigned short    trans_src_port;
    char              trans_dst_ver;
    unsigned short    trans_dst_port;
    unsigned short    policy;
    unsigned char     policy_len;
    unsigned char     action;
    unsigned char     linkgrp_len;
    unsigned char     link_len;

```

```
        unsigned char        data[];//src+dst+policy
    } tlog_llb_t;

typedef struct _tlog_l7_tcps {
    unsigned short    next_len;
    unsigned int      itime;
    unsigned int      duration;
    uint64_t          ibytes;
    uint64_t          obytes;
    unsigned char      proto;
    unsigned char      srv;
    unsigned char      src_ver;
    unsigned short     src_port;
    unsigned char      dst_ver;
    unsigned short     dst_port;
    unsigned char      trans_src_ver;
    unsigned short     trans_src_port;
    unsigned char      trans_dst_ver;
    unsigned short     trans_dst_port;
    unsigned char      action;
    unsigned short     policy;
    unsigned char      policy_len;
    unsigned char      data[];//src+dst+trans_src+trans_
dst+policy+method+host+agent+url+qry+cookie
} tlog_l7_tcps_t;

typedef struct _tlog_l7_radius {
    unsigned short    next_len;
    unsigned int      itime;
    unsigned int      duration;
    uint64_t          ibytes;
    uint64_t          obytes;
    unsigned char      proto;
    unsigned char      srv;
    unsigned char      src_ver;
    unsigned short     src_port;
    unsigned char      dst_ver;
    unsigned short     dst_port;
    unsigned char      trans_src_ver;
    unsigned short     trans_src_port;
    unsigned char      trans_dst_ver;
    unsigned short     trans_dst_port;
    unsigned char      action;
    unsigned short     policy;
    unsigned char      policy_len;
    unsigned char      user_len;
    unsigned char      data[];//src+dst+trans_src+trans_dst+policy+user
} tlog_l7_radius_t;
```

```

typedef struct _tlog_l7_dns {
    unsigned short    next_len;
    unsigned int      itime;
    unsigned int      duration;
    uint64_t          ibytes;
    uint64_t          obytes;
    unsigned char      proto;
    unsigned char      srv;
    unsigned char      src_ver;
    unsigned short     src_port;
    unsigned char      dst_ver;
    unsigned short     dst_port;
    unsigned char      action;
    unsigned short     policy;
    unsigned char      policy_len;
    unsigned char      domain_len;
    unsigned char      res_ip_ver;
    unsigned char      data[]; //src+dst+trans_src+trans_dst+policy+domain
} tlog_l7_dns_t;

typedef struct _tlog_l7_http {
    unsigned short    next_len;
    unsigned int      itime;
    unsigned int      duration;
    uint64_t          ibytes;
    uint64_t          obytes;
    unsigned char      proto;
    unsigned char      srv;
    unsigned char      src_ver;
    unsigned short     src_port;
    unsigned char      dst_ver;
    unsigned short     dst_port;
    unsigned char      trans_src_ver;
    unsigned short     trans_src_port;
    unsigned char      trans_dst_ver;
    unsigned short     trans_dst_port;
    unsigned char      action;
    unsigned short     method_len;
    unsigned short     host_len;
    unsigned short     agent_len;
    unsigned short     qry_len;
    unsigned short     cookie_len;
    unsigned short     url_len;
    unsigned short     retcode;
    unsigned short     policy;
    unsigned char      policy_len;
    unsigned char      data[]; //src+dst+trans_src+trans_
dst+policy+method+host+agent+url+qry+cookie
} tlog_l7_http_t;

```

```
typedef struct _alog_l4
{
    unsigned short    next_len;
    unsigned int      itime;
    unsigned int      duration;
    unsigned int      count;
    unsigned char     sub_severity;
    unsigned char     irdb_type;
    unsigned char     proto;
    unsigned char     src_ver;
    unsigned short    src_port;
    unsigned char     dst_ver;
    unsigned short    dst_port;
    unsigned short    policy;
    unsigned char     policy_len;
    unsigned char     action;
    unsigned char     data[]; //src+dst+policy
} alog_l4_t;

typedef struct _alog_l7 {
    unsigned short    next_len;
    unsigned int      itime;
    unsigned int      duration;
    uint64_t          ibytes;
    uint64_t          obytes;
    unsigned char     proto;
    unsigned short    srv;
    unsigned char     src_ver;
    unsigned short    src_port;
    unsigned char     dst_ver;
    unsigned short    dst_port;
    unsigned char     action;
    unsigned short    method_len;
    unsigned short    host_len;
    unsigned short    agend_len;
    unsigned short    url_len;
    unsigned short    retcode;
    unsigned short    policy;
    unsigned char     policy_len;
    unsigned char     data[]; //src+dst+method+host+agent+url+policy
} alog_l7_t;

typedef struct _alog_syn {
    unsigned short    next_len;
    unsigned int      itime;
    unsigned int      duration;
    unsigned int      count;
    unsigned char     sub_severity;
```




High Performance Network Security



Copyright© 2017 Fortinet, Inc. All rights reserved. Fortinet®, FortiGate®, FortiCare® and FortiGuard®, and certain other marks are registered trademarks of Fortinet, Inc., in the U.S. and other jurisdictions, and other Fortinet names herein may also be registered and/or common law trademarks of Fortinet. All other product or company names may be trademarks of their respective owners. Performance and other metrics contained herein were attained in internal lab tests under ideal conditions, and actual performance and other results may vary. Network variables, different network environments and other conditions may affect performance results. Nothing herein represents any binding commitment by Fortinet, and Fortinet disclaims all warranties, whether express or implied, except to the extent Fortinet enters a binding written contract, signed by Fortinet's General Counsel, with a purchaser that expressly warrants that the identified product will perform according to certain expressly-identified performance metrics and, in such event, only the specific performance metrics expressly identified in such binding written contract shall be binding on Fortinet. For absolute clarity, any such warranty will be limited to performance in the same ideal conditions as in Fortinet's internal lab tests. In no event does Fortinet make any commitment related to future deliverables, features, or development, and circumstances may change such that any forward-looking statements herein are not accurate. Fortinet disclaims in full any covenants, representations, and guarantees pursuant hereto, whether express or implied. Fortinet reserves the right to change, modify, transfer, or otherwise revise this publication without notice, and the most current version of the publication shall be applicable.